ORANSIM: SIMULATING CONSUMER RESPONSE TO SOCIAL MEDIA CAMPAIGNS BEFORE LAUNCH

Anonymous authors
Paper under double-blind review

ABSTRACT

Social simulation studies how individual behavior and social interaction produce collective outcomes. In social media marketing, campaign actions shape which consumers encounter the content and how they respond; these responses then spread through the population. We propose **OranSim**, a social simulation framework that connects creative, creator, targeting, and budget choices to this process. Heterogeneous consumers receive exposure according to content matching and platform allocation and generate initial responses, which propagate among 60 population segments. Candidate campaigns share the initial population and aligned random numbers, making their response trajectories comparable under action changes. In a controlled synthetic campaign, doubling the budget approximately doubles reach while lowering mean content match and engagement probability among the reached consumers; mean 14-day cumulative simulated response mass rises to 1.96 times the baseline. LightGBM predictors fitted to 39,000 historical RedNote notes estimate platform engagement with log-scale R^2 of 0.56–0.62 in five-fold cross-validation; a separate 12,154-note corpus supplies temporal, unseen-creator, and held-out-niche test splits. Public-data experiments evaluate policy-value estimation and audience ranking, and paired synthetic outcomes test counterfactual scoring. Together, scenario trajectories and engagement estimates support campaign selection according to a prespecified marketing objective.

1 INTRODUCTION

Social simulation studies how the behavior of individuals, their interactions, and their environment give rise to collective outcomes. Agent-based models represent heterogeneous people and their interactions to study product diffusion, social activity, and financial markets (Rand & Rust, 2011; Rand & Stummer, 2021; Piao et al., 2025; Yang et al., 2025). Social media marketing offers a concrete setting for this question: brands publish content through creators, platforms allocate exposure, and consumers' initial responses seed subsequent propagation (Hughes et al., 2019; Bakshy et al., 2015; 2012). Campaign response depends on how marketing actions enter this process.

Creative choices change content matching; creator selection changes the source and audience match; targeting changes who receives exposure; and budget determines reach. These actions interact through the audience they jointly select: creative, creator, and targeting choices shape the exposure ranking, and budget determines how many of the ranked consumers receive exposure. Expanding reach therefore admits consumers with different response tendencies, so the resulting population response depends on both the number of people reached and their composition. Seeding research similarly connects firms' choice of initial recipients to participation and subsequent diffusion (Hinz et al., 2011). Brands choose among these options before launch, when the consumers who will be reached and their responses are still unobserved. How can a simulator carry each option through exposure, individual response, and propagation from a shared initial population?

Market-response models connect marketing inputs to aggregate outcomes (Hanssens et al., 2001; Jin et al., 2017); prelaunch comparison additionally requires tracing the consumers reached by each option. Cascade models forecast propagation from post-publication events (Zhao et al., 2015; Cao et al., 2017); the prelaunch decision precedes those observations. Social simulators represent individual decisions and interaction, and world models compute action-conditioned environment evolution (Piao et al., 2025; Yang et al., 2025; Ha & Schmidhuber, 2018; Hafner et al., 2025).

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Applying this simulation perspective to marketing requires introducing creative, creator, targeting, and budget into the action space and connecting them to population dynamics.

We propose **OranSim**, a social simulation framework for social media marketing. Its central idea is to express marketing actions as changes to content matching, exposure allocation, and reach, then propagate the initial responses of the selected audience. Consumer interests, activity, personality, and response states determine exposure and behavioral probabilities, and Figure 1 marks where each attribute enters. Initial responses propagate between demographic segments. Candidate scenarios share their initial population, mechanism parameters, and individual random numbers, so action changes can be compared on the same people.

Campaign objectives are stated in platform engagement counts: reads, likes, collects (saves), and comments. Predictors learned from historical RedNote posts, called notes, estimate these counts from the note content, creator, and timing of each scenario; the objective specifies how to use them alongside the scenario's simulated trajectory. In the controlled case, a collect-first objective ranks options by predicted collects and orders options with equal predictions by cumulative simulated response. This use of historical observations alongside simulation follows a prelaunch forecasting tradition in marketing (Trusov et al., 2013).

Our contributions are:

1. A marketing social simulation that maps campaign actions to the exposure and response of heterogeneous consumers, and connects their initial responses to population dynamics through intergroup propagation.
2. Controlled scenario comparison from shared initial conditions. A skincare campaign traces how doubling the budget expands reach, changes audience composition, and raises mean cumulative simulated response mass to 1.96 times the baseline, then applies a prespecified objective to select a campaign.
3. Evaluation of the predictors and estimators used in campaign comparison: platform engagement prediction on private RedNote notes, policy-value estimation and audience ranking on public data, and counterfactual scoring on paired synthetic outcomes, with fixed scenarios and a deterministic output snapshot for reproduction.

2 Related Work

Social simulation. Generative Agents, RecAgent, OASIS, and AgentSociety model individual behavior and interaction in social and recommendation environments (Park et al., 2023; Wang et al., 2025; Yang et al., 2024; Piao et al., 2025). TwinMarket connects investor decisions, information propagation, and collective market dynamics (Yang et al., 2025). Language-model studies examine behavioral responses and human alignment (Argyle et al., 2023; Aher et al., 2023; Brand et al., 2024; Filippas et al., 2024; Ye et al., 2026; Hu et al., 2026); world models, recommendation simulators, and population frameworks support action-conditioned prediction at different scales (Ha & Schmidhuber, 2018; Hafner et al., 2025; Zhao et al., 2023; Zhang et al., 2024; 2025b;a; Gao et al., 2024). OranSim represents a campaign as an action whose consequences unfold through statistical individual responses and intergroup propagation.

Consumer interaction and market diffusion. The Bass model describes new-product adoption through innovation and imitation (Bass, 1969), and agent-based marketing links consumer heterogeneity and local interaction to aggregate diffusion (Rand & Rust, 2011; Rand & Stummer, 2021). Segmentation, persuasion, and interpersonal influence inform these relationships (McCarthy, 1960; Kotler & Keller, 2016; Petty & Cacioppo, 1986; Katz & Lazarsfeld, 1955; Krugman, 1972). Seeding strategies connect the choice of initial recipients to participation and reach (Hinz et al., 2011). For prelaunch forecasting, pretest-market models estimate new-product sales from consumer responses measured before launch (Silk & Urban, 1978), and synthetic-network priors combine simulation with historical diffusion (Trusov et al., 2013). OranSim uses campaign choices to determine the initial audience and its response, then computes propagation over that population.

Outcome prediction and decision evaluation. Market-response models estimate aggregate outcomes from marketing inputs (Hanssens et al., 2001; Jin et al., 2017). Hawkes processes and cascade

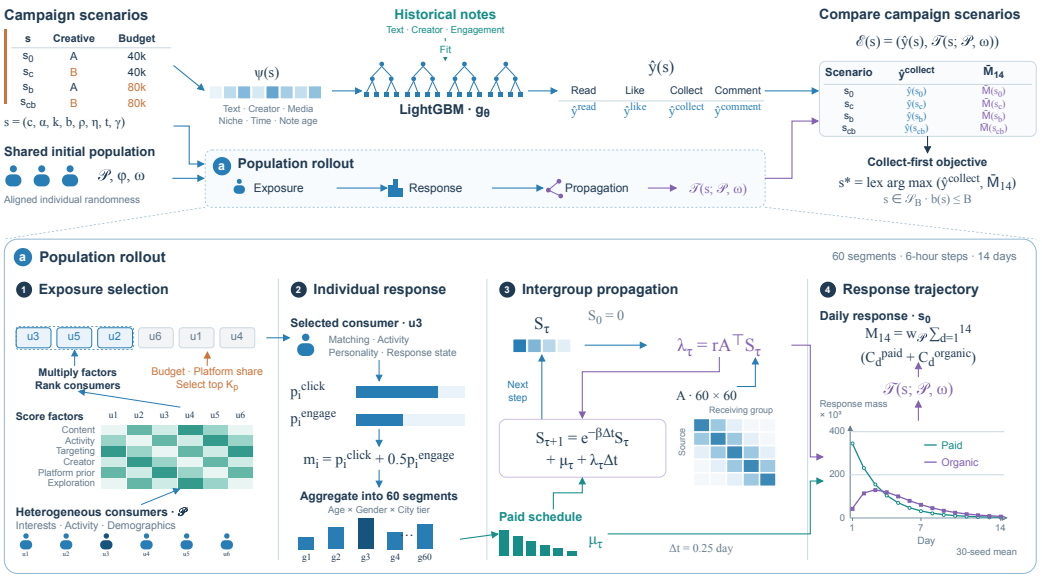

Figure 1: OranSim overview: note-engagement prediction and population rollout for each scenario (top), the four rollout stages (bottom). Curves are 30-seed means for baseline s_0 (Section 5.2).

predictors model propagation from observed events (Hawkes, 1971; Zhao et al., 2015; Li et al., 2017; Cao et al., 2017; Wang et al., 2017a). Off-policy estimators use logged assignments to estimate policy value (Li et al., 2011; Dudík et al., 2011; Swaminathan & Joachims, 2015; Wang et al., 2017b), with KuaiRand and Open Bandit providing public evaluation data (Gao et al., 2022; Saito et al., 2020). Treatment-effect learners rank audiences by incremental response (Künzel et al., 2019; Nie & Wager, 2021; Athey et al., 2019). Our evaluation uses these outcome and decision measures.

3 OranSim

OranSim computes how a campaign changes the response of a simulated society. A population of heterogeneous consumers supplies interests, activity, and behavioral attributes; exposure allocation selects the initial audience; individual responses seed propagation between demographic segments. Historical-note predictors estimate platform engagement for the same scenario, and the campaign objective specifies how to use the engagement estimates and rollout trajectory. Figure 1 shows the computation and the consumer attributes used at each stage.

3.1 Simulated Society for Marketing

The population $\mathcal{P}$ contains consumers with demographic attributes, interest vectors, Big Five personality traits, platform activity, and response states. These attributes enter distinct parts of the simulation. Interests determine content and creator matching; demographics determine targeting weights and platform audience priors; platform activity affects exposure ranking and response probability. Openness and neuroticism enter the statistical response model, and a response-state coordinate supplies a fatigue proxy. Thus two consumers who receive the same creative can have different response probabilities. Appendix A.1 specifies population generation.

Age band, gender, and city tier also define $6 \times 2 \times 5 = 60$ segments. A matrix A describes the influence of each source segment on receiving segments, assigning greater weight within a segment and between demographically similar segments. Individual responses are aggregated into these segments before propagation.

3.2 From Individual Exposure to Population Response

For consumer i on platform p, the exposure score multiplies content matching, platform activity, targeting weight, creator matching, the platform audience prior, and a seeded exploration perturbation. Consumers are selected in decreasing score order. With budget b, platform share ρ_p, population scaling weight $w_{\mathcal{P}}$, and cost per thousand impressions CPM_p, the number selected is

$$K_p = \min\left(|\mathcal{P}|, \left\lfloor \frac{1000 b \rho_p}{w_{\mathcal{P}} \mathrm{CPM}_p} \right\rfloor\right). \tag{1}$$

Here $w_{\mathcal{P}}$ is the ratio of represented to simulated population size. Increasing the budget extends selection further down the ranking, changing both reach and the composition of the exposed audience. Appendix A.2 defines each factor, including the hashed title representation for content matching.

For each selected consumer, a logistic behavior model computes click probability p_i^{click} from matching, activity, targeting, creator, personality, and state features. Multiplying this probability by a logistic post-click engagement propensity gives p_i^{engage}. The initial response mass $p_i^{\mathrm{click}} + 0.5 p_i^{\mathrm{engage}}$ counts expected clicks and weights expected engagement by one half. We sum these masses within each segment and distribute them over time according to the paid schedule, obtaining the segment vector $\boldsymbol{\mu}_\tau$ at step τ. The model coefficients and schedule are specified in Appendix A.3.

A segmented Hawkes process (Hawkes, 1971) propagates these initial responses. Let $\mathbf{S}_\tau$ be accumulated excitation at the start of a step, and let A_{ij} be the influence of source segment i on segment j. The organic response rate and next state are

$$\boldsymbol{\lambda}_\tau = r A^\top \mathbf{S}_\tau, \qquad \mathbf{S}_{\tau+1} = e^{-\beta \Delta t} \mathbf{S}_\tau + \boldsymbol{\mu}_\tau + \boldsymbol{\lambda}_\tau \Delta t, \quad \mathbf{S}_0 = \mathbf{0}, \tag{2}$$

where r is the branching parameter that scales propagation excitation, β is the decay rate, and Δt is measured in days. Each step combines retained excitation, current paid input, and organic response induced by earlier activity. Summing $\boldsymbol{\mu}_\tau$ and $\boldsymbol{\lambda}_\tau \Delta t$ over segments and within each day gives daily paid and organic response masses C_d^{paid} and C_d^{organic}. Aggregating the steps within each day separately for each segment gives its response trajectory.

3.3 Scenario Rollouts of Marketing Actions

A campaign scenario is

$$s = (c, \alpha, k, b, \boldsymbol{\rho}, \eta, t, \gamma), \tag{3}$$

where c contains the title, body, media, niche, and hashtags; α specifies the audience; k describes the creator; b is the budget; and $\boldsymbol{\rho}$ contains nonnegative platform budget shares summing to one. The exposure configuration η gives platform priors, exploration strength, costs, and the paid schedule. Publication time is t; $\gamma \in [0, 180]$ is the note age in days at which outcomes are predicted.

Each action has a concrete point of entry. Creative c changes content matching in exposure and response; creator k supplies niche and audience-interest matching. Audience setting α changes demographic targeting weights. Budget b and shares $\boldsymbol{\rho}$ determine reach through Equation (1). For example, replacing a skincare creative recomputes content matching, the selected audience, and its responses while holding the budget and initial population fixed.

The rollout $\mathcal{T}(s; \mathcal{P}, \omega) = R_\phi(s, \mathcal{P}; \omega)$ performs exposure, response, and propagation in sequence, where ϕ contains their mechanism parameters and ω is a random seed. Candidate scenarios share $\mathcal{P}$ and ϕ, and each individual uses aligned random numbers across scenarios. Exposure and propagation states are computed separately for each candidate. This pairing reduces variation due to population resampling when comparing action changes.

The cumulative simulated response over H days is

$$M_H(s, \omega) = w_{\mathcal{P}} \sum_{d=1}^{H} \big(C_d^{\mathrm{paid}}(s, \omega) + C_d^{\mathrm{organic}}(s, \omega) \big). \tag{4}$$

It measures weighted response mass in the represented population. For baseline s_0 and candidate s_1, the controlled difference is

$$\Delta_{\mathrm{sim}}(s_1, s_0; \omega) = M_{14}(s_1, \omega) - M_{14}(s_0, \omega). \tag{5}$$

We average this difference over common seeds. Daily trajectories retain the timing of response, while M_{14} compares the total over the same window.

3.4 Observed Engagement Prediction

Historical notes pair the content, creator, and timing of published notes with their engagement counts. For each note, the predictor concatenates principal components of title and body representations with log follower count, duration, encoded media type and niche, hashtag features, cyclic publication-time features, and normalized note age $\gamma/180$; for training notes, age counts the days from publication to the fitting date, capped at 180. Denote this scenario-to-feature map by $\psi(s)$. We compare configurations without and with note age, called v3 and v3.1-pg, respectively. Each trains four separate LightGBM regressors (Ke et al., 2017) on $\log(1+y)$, one each for reads, likes, collects, and comments, and reports $\exp(\hat{z})-1$ as the predicted count for each log-scale prediction $\hat{z}$. Separate outcomes allow a campaign to prioritize its chosen engagement objective.

3.5 Prelaunch Comparison

For a scenario, the evaluation record combines its engagement predictions and trajectory:

$$\mathcal{E}(s;\mathcal{P},\omega) = \big(\hat{\mathbf{y}}(s), \mathcal{T}(s;\mathcal{P},\omega)\big) = \big(g_\theta(\psi(s)), R_\phi(s,\mathcal{P};\omega)\big), \tag{6}$$

where g_θ denotes the four fitted regressors. Both computations receive the same scenario and describe different objects: $\hat{\mathbf{y}}(s)$ predicts the engagement of the campaign note, and the rollout traces the consumers the campaign reaches and their responses. Audience, budget, platform shares, and exposure configuration enter only the rollout, where they shape who is reached; a budget-only change therefore leaves $\psi(s)$ unchanged and selects more consumers. The campaign objective determines how to compare the resulting quantities.

For the collect-first objective with budget cap B, let $\mathcal{S}_B = \{s \in \mathcal{S} : b(s) \le B\}$. We select

$$s^\star = \operatorname*{lex\,arg\,max}_{s\in\mathcal{S}_B} \big(\hat{y}^{\text{collect}}(s), \overline{M}_{14}(s)\big), \tag{7}$$

where lexicographic maximization first selects the highest predicted collect count and then resolves ties by mean cumulative simulated response. Other objectives order the quantities in $\mathcal{E}$ differently: an objective centered on population response ranks candidates by $\overline{M}_{14}$ or by response per CNY, $\overline{M}_{14}/b$.

4 Evaluation Protocol

Each evaluation examines one computation used in campaign comparison with data that can measure it: observed engagement of private RedNote notes tests the note predictors g_θ; the controlled campaign traces how actions pass through the rollout R_ϕ; public logs with randomized exposure, recorded assignment probabilities, or randomized treatment test policy-value estimators for content allocation and treatment-effect learners for audience ranking; and paired synthetic outcomes test feature-substitution scoring.

4.1 Prediction and Controlled Response

We report the coefficient of determination $R^2_{\log}$ separately for reads, likes, collects, and comments after applying $\log(1+y)$ to observed and predicted counts. Shared folds compare predictor configurations; structured holdouts test temporal, creator, and niche generalization. Controlled-campaign rollouts record reach, mean content matching, engagement probability, and paid and organic response.

4.2 Off-Policy Value Estimation

A policy $\pi(a \mid x)$ assigns content or actions a in context x; its value is the expected outcome under that assignment. Given logged outcome y_i and logging probability μ_i, define $w_i = \pi(a_i \mid x_i)/\mu_i$. Inverse propensity scoring (IPS) estimates the value by $\frac{1}{n}\sum_i w_i y_i$ and self-normalized IPS (SNIPS) by $\sum_i w_i y_i / \sum_i w_i$; the effective sample size $n_{\text{eff}} = (\sum_i w_i)^2 / \sum_i w_i^2$ describes weight concentration. Doubly robust (DR) estimation combines outcome predictions and weighted residuals; Switch-DR applies the residual correction below a weight threshold (Li et al., 2011; Swaminathan & Joachims, 2015; Dudík et al., 2011; Wang et al., 2017b). Appendix D.1 specifies these estimators and their identification conditions.

Table 1: RedNote prediction measured by $R^2_{\log}$. The first two rows share notes and folds. Structured holdouts refit predictors on a separate corpus with hashed text features.

Evaluation setting	Evaluation size	Reads	Likes	Collects	Comments
Five-fold CV, without age (v3)	39,000	0.5752	0.6122	0.5766	0.5492
Five-fold CV, with age (v3.1-pg)	39,000	**0.5928**	**0.6223**	**0.5869**	**0.5585**
Latest-time holdout	1,823 test notes	0.253	0.361	0.329	-0.007
Unseen-creator holdout	1,844 test notes	0.559	0.550	0.513	0.347
Leave-one-niche-out	12,154 / 13 folds	0.503	0.513	0.456	0.333

4.3 Audience Ranking and Paired Outcomes

For advertiser-randomized interventions, the conditional average treatment effect (CATE) is $\tau(x) = \mathbb{E}[Y(1) - Y(0) \mid X = x]$, where $Y(1)$ and $Y(0)$ are outcomes with and without intervention. We rank customers by estimated incremental response using treatment-effect learners (Künzel et al., 2019; Nie & Wager, 2021; Athey et al., 2019). For the top fraction q, uplift@q is the treatment-minus-control difference in mean observed outcomes, $\bar{y}(T_q^{(1)}) - \bar{y}(T_q^{(0)})$, where $T_q^{(1)}$ and $T_q^{(0)}$ are the treatment and control customers in the selected set. Normalized area under the uplift curve (AUUC) and Qini summarize performance across coverage fractions (Appendix D.2).

The paired synthetic test evaluates factual fit, alternative-outcome error, effect direction, and grouped CATE fit for fixed predictors under feature substitution (Section 5.4).

5 Experiments

OranSim uses four LightGBM regressors for note engagement, a logistic response model for individual click and engagement probabilities, and a 60-segment Hawkes process for propagation. The prediction experiments fit the engagement regressors; the campaign case holds the fitted predictor and response mechanisms fixed while changing actions. Public-data experiments use off-policy estimators and treatment-effect learners for content allocation and audience selection. The paired synthetic experiment fits four more LightGBM regressors to score alternatives with known outcomes.

5.1 RedNote Engagement Prediction

The five-fold comparison uses 39,000 private RedNote note snapshots and fits a LightGBM regressor for each of read, like, collect, and comment counts, with $\log(1 + y)$ targets. The configurations without and with note age, v3 and v3.1-pg, share shuffled folds with seed 42. Titles and bodies are embedded separately with OpenAI text-embedding-3-small, each truncated to 768 dimensions; the concatenated 1,536 dimensions are reduced to 128 by PCA; adding metadata gives 182 features without note age and 183 with it.

The structured holdouts use a separately assembled sample of 12,154 notes from 13 niches, each with at least 100 reads. Each split refits 300-tree LightGBM regressors using 768-dimensional deterministic hashed representations of titles and bodies, concatenated and reduced to at most 64 dimensions by training-set PCA. The temporal split uses 10,331/1,823 training/test notes; the creator split uses 10,310/1,844, with 1,642 test creators absent from training; leave-one-niche-out retrains on 13 folds. Appendix C gives regressor settings, features, splits, and errors.

Table 1 shows that note age improves all four outcomes: the gains in unrounded $R^2_{\log}$ are 0.0176, 0.0101, 0.0103, and 0.0092 for reads, likes, collects, and comments. The resulting scores of 0.5585–0.6223 support engagement prediction for comparing candidate options. The unseen-creator scores of 0.347–0.559 and leave-one-niche-out scores of 0.333–0.513 support prediction for creators and niches absent from training. The latest-time scores are 0.253, 0.361, 0.329, and -0.007, respectively.

Table 2: Controlled campaign options. Inputs give creative, budget in CNY, and creator tier. Reads, likes, collects, and comments are predicted note counts at age 14 days; $\overline{M}_{14}$ is mean cumulative simulated response mass over 30 seeds.

Option	Inputs	Reads	Likes	Collects	Comments	$\overline{M}_{14}$
s_0	A / 40,000 / mid	54,819	2,209	1,371	52	1,834,405
s_c	B / 40,000 / mid	56,520	2,016	1,020	80	1,832,978
s_b	A / 80,000 / mid	54,819	2,209	1,371	52	3,599,734
s_k	A / 40,000 / high	60,882	2,672	1,480	79	1,833,903
s_{cb}	B / 80,000 / mid	56,520	2,016	1,020	80	3,604,845

Table 3: The response process behind Table 2. Reach counts sampled individuals; match and engagement are means among those reached; $\overline{M}_{14}$ is scaled to the represented population. The last two columns give the percentage by which $\overline{M}_{14}$ exceeds two equal-reach controls, uniform selection and mean individual features. Means over the 30 shared seeds; Appendix B.2 adds standard deviations.

Option	Sample reach	Content match	Engagement prob.	$\overline{M}_{14}$	vs. uniform	vs. mean
s_0	41,666	0.5009	0.7020	1,834,405	+6.61%	−0.17%
s_c	41,666	0.4971	0.7009	1,832,978	+6.19%	−0.16%
s_b	83,333	0.4942	0.6783	3,599,734	+4.59%	−0.19%
s_k	41,666	0.5011	0.7015	1,833,903	+6.59%	−0.17%
s_{cb}	83,333	0.4987	0.6799	3,604,845	+4.42%	−0.17%

5.2 Simulated Social Response to Marketing Actions

The controlled skincare campaign follows an action change through exposure, individual response, and propagation. We sample 100,000 individuals representing 2,000,000 people, fix the population with seed 2027, and compare five options with 30 shared seeds $\{0, \ldots, 29\}$. Note predictions use the frozen v3.1-pg LightGBM predictor with its training text encoder at an evaluation age of 14 days. Response parameters stay fixed, and propagation uses $\beta = 0.9$ per day, $r = 0.35$, and six-hour steps (Appendix A).

The baseline combines informational creative A, a mid-tier creator, and CNY 40,000. Creative B presents the same triple-ceramide repair routine as a response to seasonal redness and invites viewers to take a skin-type test and claim a trial-size sample. Both are 30-second videos. Candidate options change the creative, budget, creator tier, or creative and budget together (Table 2). Targeting, platform, publication time, and exposure configuration are shared; Appendix B gives the full campaign.

Increasing the budget expands reach and changes the composition of the reached audience. With creative A, doubling the budget selects 83,333 individuals in decreasing exposure-score order, up from 41,666. Mean content match falls from 0.5009 to 0.4942, and mean engagement probability from 0.7020 to 0.6783 (Table 3). Their initial responses enter the paid schedule and excite subsequent organic responses. Under the shared propagation parameters, $\overline{M}_{14}$ rises from 1,834,405 to 3,599,734: reach grows to 2.00 times the baseline and response mass to 1.96 times. The added individuals rank lower in exposure score and have lower mean response propensity, so response mass grows slightly less than reach.

Two equal-reach controls locate these differences (Table 3). Uniform selection of the same number of individuals lowers the mean engagement probability for s_0 from 0.7020 to 0.6346, and the full rollout exceeds it by 4.4%–6.6% in $\overline{M}_{14}$; giving every selected individual the audience's mean features changes $\overline{M}_{14}$ by less than 0.2%. Heterogeneity thus shapes aggregate response mainly through which consumers the exposure ranking selects. With fixed initial segment responses, uniform cross-segment weights lower the targeted segments' organic share from 96.1% under A to 94.3%, with total response unchanged (Appendix B.6).

At CNY 40,000, replacing A with B preserves the reach of 41,666 individuals and changes mean match from 0.5009 to 0.4971, engagement probability from 0.7020 to 0.7009, and response mass from 1,834,405 to 1,832,978, a paired difference of $-1{,}427$ (95% bootstrap interval $[-1{,}510, -1{,}343]$). The historical-note predictor gives A 1,371 collects and B 1,020; B has more predicted reads and

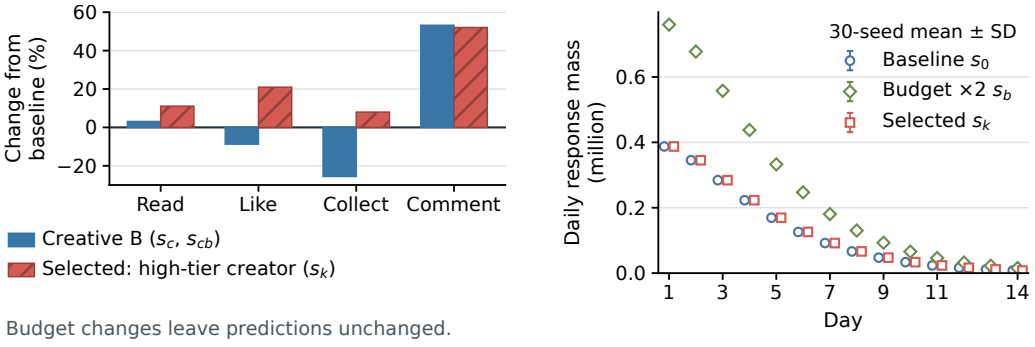

(a) Predicted changes from the baseline.

(b) Daily response of three options.

Figure 2: Campaign comparison in the synthetic skincare setting. (a) Changes in four note predictions; budget changes leave them unchanged, so s_b matches the baseline and s_{cb} matches creative B. (b) Daily simulated response, mean ± 1 standard deviation over 30 seeds; the nearly coinciding s_0 and s_k points are offset horizontally. The collect-first rule selects the high-tier creator at CNY 40,000.

comments and fewer likes. The high-tier creator raises all four predictions, to 1,480 collects and 79 comments.

Under the CNY 80,000 cap, the collect-first rule in Equation (7) selects s_k, which has the most predicted collects. Options that differ only in budget share their predictions, and cumulative response mass orders each such pair, placing s_b ahead of s_0 and s_{cb} ahead of s_c. The comparison thus connects the engagement objective to creator and creative choice and the population response to budget choice. Under objectives centered on population response, the rollout decides: ranking by $\overline{M}_{14}$ selects s_{cb} and ranking by response per CNY selects s_0, as doubling the budget raises response mass by 96.2% and lowers response per CNY by 1.9%. All three rankings hold under eleven single-parameter changes to propagation, targeting, exploration, and response parameters, which keep the budget-doubling ratio within 1.950–1.973 (Appendix B.7). Figure 2 shows the predictions and daily responses. Repeated same-seed baselines have zero aggregate and daily differences (Appendix B.5).

5.3 Policy Value and Audience Ranking on Public Data

We evaluate content allocation as policy-value estimation and audience selection as treatment-effect ranking. KuaiRand-Pure supplies 1,186,059 randomized exposure records from 27,285 users and 7,583 videos, with logging probability $1/7{,}583$ and long view as the outcome (Gao et al., 2022). Open Bandit supplies 10,000 records each from random and Bernoulli Thompson Sampling (BTS) policies over 80 actions (Saito et al., 2020). Each log is split in time: the earlier part selects the ten most-logged actions, shared by the two Open Bandit logs, and fits the reward means used by DR and Switch-DR; the later part is evaluated. The target policy allocates probability 0.1 uniformly over all actions and 0.9 uniformly over the ten. Intervals come from 1,000 bootstrap resamples of users on KuaiRand and of timestamps on Open Bandit (Appendix D.1).

X5 RetailHero contains 200,039 customers, treatment fraction 0.4998, and outcome rate 0.6199. Seeds $\{42, 137, 256\}$ repeat training with shared 75%/25% treatment/outcome-stratified splits: 150,029 training and 50,010 test customers per seed. The methods are random ranking, two-model logistic regression, S/T/X learners and DRLearner with 80-tree forests, and CausalForestDML with 200 causal trees; Appendix D.3 gives configurations and every reach fraction.

On KuaiRand, SNIPS is 0.0879 with interval $[0.0755, 0.1035]$, and IPS, DR, and Switch-DR lie within 0.0828–0.0932. The target policy concentrates probability on ten videos, so their records carry weight 682.57 and all others 0.10; these heavily weighted records reduce the effective sample size to 1,454 of 847,973 records. The mean weight is 1.060 with interval $[1.005, 1.114]$. On the random and BTS Open Bandit logs, with 19 and 17 evaluated clicks, SNIPS is 0.0100 with interval $[0.0034, 0.0184]$ and 0.0022 with $[0.0007, 0.0040]$ (Table D.1).

Table 4: LightGBM feature-substitution scoring on 2,000 paired synthetic test scenarios. Counterfactual normalized MAE (nMAE) divides MAE by the mean absolute alternative outcome.

Outcome	Factual R^2	Counterfactual nMAE	Effect-sign accuracy	Grouped CATE R^2
Exposure	0.964	0.219	0.955	0.457
Clicks	0.892	0.256	0.617	0.355
Conversions	0.762	0.582	0.930	0.531
Revenue	0.667	0.439	0.552	0.003

On X5, two-model logistic regression has the highest mean AUUC and X-Learner the highest Qini; DRLearner leads in uplift at 5% and 10% reach and X-Learner at 20% and 30%, so the preferred ranking method depends on the reachable fraction of customers (Appendix D.3).

5.4 Paired Counterfactual Scoring

We test feature-substitution scoring using 20,000 synthetic scenarios with known factual and alternative outcomes, split by scenario identifier into 16,000/2,000/2,000 training/validation/test cases. The alternative setting changes the budget by a prespecified rule while preserving creator attributes. The generator computes exposure, clicks, conversions, and revenue under each setting (Appendix E.1).

Four 300-tree LightGBM regressors fit the factual training outcomes. Each test scenario is scored twice with fixed predictor parameters, substituting the factual and alternative features; their difference estimates the intervention effect (Table 4). CATE evaluation averages generated and predicted differences within budget tiers, creator tiers, and niches before computing R^2; within-group averaging reduces the influence of the independently drawn outcome perturbations. Appendix E gives all metrics.

6 Limitations

The outcome predictors fit a private observational RedNote corpus selected by a minimum engagement threshold and within-niche engagement rankings. They model the platform distribution represented by this corpus, and their prediction differences between creatives or creators reflect associations observed in these notes. Back-transformed counts fall 57%–82% below observed means in the structured holdouts, and the collect-first rule uses only their order (Appendix C.3). Training note age is counted to the fitting date rather than to each count's collection, which changes temporal errors (Appendix C.4). Retraining the predictors requires the private corpus and the text-embedding-3-small service; the deterministic snapshot checks scoring with hashed text features.

The four note-outcome predictors are fitted to observations; response, propagation, and intergroup influence use design parameters, which also set absolute response levels, so the case compares options by differences. Calibrating these parameters and testing intervention differences and trajectory shapes on real platforms require campaign-level observations. Rollout content matching uses hashed titles without semantic relations to interests. X5 supplies randomized treatment labels and observed responses for ranking evaluation; individual conditional treatment effects remain unobservable.

7 Conclusion

OranSim connects marketing actions to population response through exposure allocation, heterogeneous individual behavior, and intergroup propagation. Shared initial populations and aligned random numbers allow candidate campaigns to be compared along this process. The controlled campaign shows how expanding reach changes audience composition and cumulative response; historical-note predictors supply engagement estimates for objective-specific selection. The prediction, public-data, and paired synthetic experiments evaluate the predictors and estimators used in these comparisons. Together, this representation and evaluation support prelaunch analysis of how a campaign reaches consumers and how their responses accumulate over time.

AI USE STATEMENT

The research ideas, paper outline, and key arguments were developed by the human authors. OpenAI Codex assisted with refining the manuscript's organization and wording and with the English translation.

REPRODUCIBILITY STATEMENT

Appendix A specifies the population, exposure and response equations, propagation parameters, and rollout procedure. Appendix B records the campaign settings, shared-seed results, mechanism controls, and parameter sensitivity. Appendices C–E describe the prediction splits, count-scale errors, note-age reference, policy estimators, ranking metrics, and synthetic experiment. Appendix F.1 documents the deterministic scoring snapshot, and Appendix F.2 lists the code and results in the supplementary material; the public-data and paired synthetic experiments can be rerun from public data and the supplied generator and scripts.

REFERENCES

Gati V. Aher, Rosa I. Arriaga, and Adam Tauman Kalai. Using large language models to simulate multiple humans and replicate human subject studies. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 337–371. PMLR, 2023. URL https://proceedings.mlr.press/v202/aher23a.html.

Lisa P. Argyle, Ethan C. Busby, Nancy Fulda, Joshua R. Gubler, Christopher Rytting, and David Wingate. Out of one, many: Using language models to simulate human samples. *Political Analysis*, 31(3):337–351, 2023.

Susan Athey, Julie Tibshirani, and Stefan Wager. Generalized random forests. *The Annals of Statistics*, 47(2):1148–1178, 2019.

Eytan Bakshy, Itamar Rosenn, Cameron Marlow, and Lada Adamic. The role of social networks in information diffusion. In *Proceedings of the 21st International Conference on World Wide Web*, pp. 519–528, 2012. doi: 10.1145/2187836.2187907. URL https://doi.org/10.1145/2187836.2187907.

Eytan Bakshy, Solomon Messing, and Lada A. Adamic. Exposure to ideologically diverse news and opinion on Facebook. *Science*, 348(6239):1130–1132, 2015. doi: 10.1126/science.aaa1160. URL https://doi.org/10.1126/science.aaa1160.

Frank M. Bass. A new product growth for model consumer durables. *Management Science*, 15(5): 215–227, 1969. doi: 10.1287/mnsc.15.5.215.

James Brand, Ayelet Israeli, and Donald Ngwe. Using llms for market research. In *Proceedings of the 25th ACM Conference on Economics and Computation (EC)*, 2024.

Qi Cao, Huawei Shen, Keting Cen, Wentao Ouyang, and Xueqi Cheng. DeepHawkes: Bridging the gap between prediction and understanding of information cascades. In *CIKM*, 2017.

Miroslav Dudík, John Langford, and Lihong Li. Doubly robust policy evaluation and learning. In *ICML*, 2011.

Apostolos Filippas, John J. Horton, and Benjamin S. Manning. Large language models as simulated economic agents: What can we learn from homo silicus? In *Proceedings of the 25th ACM Conference on Economics and Computation (EC)*, pp. 614–615. Association for Computing Machinery, 2024. doi: 10.1145/3670865.3673513.

Chen Gao, Xiaochong Lan, Nian Li, Yuan Yuan, Jingtao Ding, Zhilun Zhou, Fengli Xu, and Yong Li. Large language models empowered agent-based modeling and simulation: A survey and perspectives. *Humanities and Social Sciences Communications*, 11, 2024. doi: 10.1057/s41599-024-03611-3.

Chongming Gao, Shijun Li, Yuan Zhang, Jiawei Chen, Biao Li, Wenqiang Lei, Peng Jiang, and Xiangnan He. KuaiRand: An unbiased sequential recommendation dataset with randomly exposed videos. In *CIKM*, pp. 3953–3957, 2022. doi: 10.1145/3511808.3557624.

David Ha and Jürgen Schmidhuber. World models. *arXiv preprint arXiv:1803.10122*, 2018. URL https://worldmodels.github.io/.

Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. Mastering diverse control tasks through world models. *Nature*, 640:647–653, 2025. doi: 10.1038/s41586-025-08744-2. URL https://www.nature.com/articles/s41586-025-08744-2.

Dominique M. Hanssens, Leonard J. Parsons, and Randall L. Schultz. *Market Response Models*. Springer, 2nd edition, 2001.

Alan G. Hawkes. Spectra of some self-exciting and mutually exciting point processes. *Biometrika*, 58(1):83–90, 1971.

Oliver Hinz, Bernd Skiera, Christian Barrot, and Jan U. Becker. Seeding strategies for viral marketing: An empirical comparison. *Journal of Marketing*, 75(6):55–71, 2011. doi: 10.1509/jm.10.0088. URL https://journals.sagepub.com/doi/10.1509/jm.10.0088.

Tiancheng Hu, Joachim Baumann, Lorenzo Lupo, Nigel Collier, Dirk Hovy, and Paul Röttger. Simbench: Benchmarking the ability of large language models to simulate human behaviors. In *International Conference on Learning Representations (ICLR)*, 2026. arXiv:2510.17516.

Christian Hughes, Vanitha Swaminathan, and Gillian Brooks. Driving brand engagement through online social influencers: An empirical investigation of sponsored blogging campaigns. *Journal of Marketing*, 83(5):78–96, 2019. doi: 10.1177/0022242919854374. URL https://doi.org/10.1177/0022242919854374.

Yuxue Jin, Yueqing Wang, Yuxiang Sun, David Chan, and Jim Koehler. Bayesian methods for media mix modeling with carryover and shape effects. Technical report, Google, 2017.

Elihu Katz and Paul F. Lazarsfeld. *Personal Influence: The Part Played by People in the Flow of Mass Communications*. Free Press, Glencoe, IL, 1955.

Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. LightGBM: A highly efficient gradient boosting decision tree. In *NeurIPS*, 2017.

Philip Kotler and Kevin Lane Keller. *Marketing Management*. Pearson, Boston, 15th edition, 2016.

Herbert E. Krugman. Why three exposures may be enough. *Journal of Advertising Research*, 12(6): 11–14, 1972.

Sören R. Künzel, Jasjeet S. Sekhon, Peter J. Bickel, and Bin Yu. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the National Academy of Sciences*, 116(10):4156–4165, 2019.

Cheng Li, Jiaqi Ma, Xiaoxiao Guo, and Qiaozhu Mei. DeepCas: An end-to-end predictor of information cascades. In *WWW*, 2017.

Lihong Li, Wei Chu, John Langford, and Xuanhui Wang. Unbiased offline evaluation of contextual-bandit-based news article recommendation algorithms. In *WSDM*, 2011.

E. Jerome McCarthy. *Basic Marketing: A Managerial Approach*. Richard D. Irwin, Homewood, IL, 1960.

Xinkun Nie and Stefan Wager. Quasi-oracle estimation of heterogeneous treatment effects. *Biometrika*, 108(2):299–319, 2021.

Joon Sung Park, Joseph C. O'Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. Generative agents: Interactive simulacra of human behavior. In *UIST*, 2023.

Richard E. Petty and John T. Cacioppo. *Communication and Persuasion: Central and Peripheral Routes to Attitude Change*. Springer, New York, 1986. doi: 10.1007/978-1-4612-4964-1.

Jinghua Piao, Yuwei Yan, Jun Zhang, Nian Li, Junbo Yan, Xiaochong Lan, Zhihong Lu, Zhiheng Zheng, Jing Yi Wang, Di Zhou, Chen Gao, Fengli Xu, Fang Zhang, Ke Rong, Jun Su, and Yong Li. AgentSociety: Large-scale simulation of LLM-driven generative agents advances understanding of human behaviors and society. *arXiv:2502.08691*, 2025. doi: 10.48550/arXiv.2502.08691. URL https://arxiv.org/abs/2502.08691.

William Rand and Roland T. Rust. Agent-based modeling in marketing: Guidelines for rigor. *International Journal of Research in Marketing*, 28(3):181–193, 2011. doi: 10.1016/j.ijresmar.2011.04.002. URL https://marketing.wharton.upenn.edu/wp-content/uploads/2016/10/Rand-W-Agent-based-modeling-in-marketing-Paper2.pdf.

William Rand and Christian Stummer. Agent-based modeling of new product market diffusion: An overview of strengths and criticisms. *Annals of Operations Research*, 305:425–447, 2021. doi: 10.1007/s10479-021-03944-1. URL https://link.springer.com/article/10.1007/s10479-021-03944-1.

Christophe Renaudin and Matthieu Martin. About evaluation metrics for contextual uplift modeling. *arXiv preprint arXiv:2107.00537*, 2021.

Yuta Saito, Shunsuke Aihara, Megumi Matsutani, and Yusuke Narita. Open Bandit Dataset and Pipeline: Towards realistic and reproducible off-policy evaluation. *arXiv:2008.07146*, 2020.

Alvin J. Silk and Glen L. Urban. Pre-test-market evaluation of new packaged goods: A model and measurement methodology. *Journal of Marketing Research*, 15(2):171–191, 1978.

Patrick D. Surry and Nicholas J. Radcliffe. Quality measures for uplift models. Technical report, 2011. URL https://www.stochasticsolutions.co.uk/pdf/kdd2011late.pdf.

Adith Swaminathan and Thorsten Joachims. The self-normalized estimator for counterfactual learning. In *NeurIPS*, 2015.

Michael Trusov, William Rand, and Yogesh V. Joshi. Improving prelaunch diffusion forecasts: Using synthetic networks as simulated priors. *Journal of Marketing Research*, 50(6):675–690, 2013. doi: 10.1509/jmr.11.0508. URL https://journals.sagepub.com/doi/10.1509/jmr.11.0508.

Lei Wang, Jingsen Zhang, Hao Yang, Zhiyuan Chen, Jiakai Tang, Zeyu Zhang, Xu Chen, Yankai Lin, Hao Sun, Ruihua Song, Xin Zhao, Jun Xu, Zhicheng Dou, Jun Wang, and Ji-Rong Wen. User behavior simulation with large language model-based agents. *ACM Transactions on Information Systems*, 43(2):55:1–55:37, 2025. doi: 10.1145/3708985. URL https://doi.org/10.1145/3708985.

Yongqing Wang, Huawei Shen, Shenghua Liu, Jinhua Gao, and Xueqi Cheng. Cascade dynamics modeling with attention-based recurrent neural network. In *IJCAI*, 2017a.

Yu-Xiang Wang, Alekh Agarwal, and Miroslav Dudík. Optimal and adaptive off-policy evaluation in contextual bandits. In *ICML*, 2017b.

Yuzhe Yang, Yifei Zhang, Minghao Wu, Kaidi Zhang, Yunmiao Zhang, Honghai Yu, Yan Hu, and Benyou Wang. TwinMarket: A scalable behavioral and social simulation for financial markets. In *Advances in Neural Information Processing Systems*, volume 38, 2025. doi: 10.52202/085713-2132. URL https://proceedings.neurips.cc/paper_files/paper/2025/hash/5bf234ecf83cd77bc5b77a24ba9338b0-Abstract-Conference.html.

Ziyi Yang, Zaibin Zhang, Zirui Zheng, Yuxian Jiang, Ziyue Gan, Zhiyu Wang, Zijian Ling, Jinsong Chen, Martz Ma, Bowen Dong, Prateek Gupta, Shuyue Hu, Zhenfei Yin, Guohao Li, Xu Jia, Lijun Wang, Bernard Ghanem, Huchuan Lu, Chaochao Lu, Wanli Ouyang, Yu Qiao, Philip Torr, and Jing Shao. OASIS: Open agent social interaction simulations with one million agents. *arXiv:2411.11581*, 2024. doi: 10.48550/arXiv.2411.11581. URL https://arxiv.org/abs/2411.11581.

Jinyi Ye, Lei Cao, Ding Chen, and Emilio Ferrara. Stop drawing scientific claims from LLM social simulations without robustness audits, 2026. URL https://arxiv.org/abs/2605.18890. arXiv:2605.18890.

An Zhang, Yuxin Chen, Leheng Sheng, Xiang Wang, and Tat-Seng Chua. On generative agents in recommendation. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 1807–1817. ACM, 2024. doi: 10.1145/3626772.3657844.

Xinnong Zhang, Jiayu Lin, Xinyi Mou, Shiyue Yang, Xiawei Liu, Libo Sun, Hanjia Lyu, Yihang Yang, Weihong Qi, Yue Chen, Guanying Li, Ling Yan, Yao Hu, Siming Chen, Yu Wang, Xuanjing Huang, Jiebo Luo, Shiping Tang, Libo Wu, Baohua Zhou, and Zhongyu Wei. SocioVerse: A world model for social simulation powered by LLM agents and a pool of 10 million real-world users, 2025a. URL https://arxiv.org/abs/2504.10157. arXiv:2504.10157.

Zijian Zhang, Shuchang Liu, Ziru Liu, Rui Zhong, Qingpeng Cai, Xiangyu Zhao, Chunxu Zhang, Qidong Liu, and Peng Jiang. LLM-powered user simulator for recommender system. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 13339–13347, 2025b. doi: 10.1609/aaai.v39i12.33456.

Kesen Zhao, Shuchang Liu, Qingpeng Cai, Xiangyu Zhao, Ziru Liu, Dong Zheng, Peng Jiang, and Kun Gai. KuaiSim: A comprehensive simulator for recommender systems. In *Advances in Neural Information Processing Systems*, volume 36, pp. 44880–44897. Curran Associates, Inc., 2023. doi: 10.52202/075280-1945. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/8c7f8f98f9a8f5650922dd4545254f28-Paper-Datasets_and_Benchmarks.pdf.

Qingyuan Zhao, Murat A. Erdogdu, Hera Y. He, Anand Rajaraman, and Jure Leskovec. SEISMIC: A self-exciting point process model for predicting tweet popularity. In *KDD*, 2015.

A Simulator Specification

This appendix gives the parameters and steps of the computation described in Sections 3.1–3.3. Table A.1 collects all constants for reference.

A.1 Simulated Population and Segments

The simulated population contains 100,000 sampled individuals representing 2,000,000 people, so the population scaling weight is $w_{\mathcal{P}} = 20$. The population is generated once with seed 2027 and shared by all candidate scenarios. Each individual has an age band, gender, city tier, income decile, education level, and occupation; a 64-dimensional unit interest vector; Big Five trait scores in $[0, 1]$; an activity level on each platform; and a 16-dimensional response state. Age band, gender, and city tier are drawn from marginal proportions, occupation is drawn conditional on age band, and income is correlated with city tier and education.

The combinations of six age bands (15–24, 25–34, 35–44, 45–54, 55–64, and 65 or older), two genders, and five city tiers (tier 1 through tier 5 and below) define the $6 \times 2 \times 5 = 60$ population segments. Exposure and individual response are computed for individuals; propagation is computed between segments.

A.2 Exposure Scoring and Budget Conversion

On platform p, the exposure score of individual i is the product of six factors,

$$e_i = m_i^{\text{content}}\, a_{ip}\, g_i^{\text{target}}\, g_i^{\text{creator}}\, \pi_{ip}\, \xi_i. \tag{A.1}$$

Content matching $m_i^{\text{content}} = (\langle \mathbf{u}_i, \mathbf{c} \rangle + 1)/2$ maps the inner product of the individual's interest vector $\mathbf{u}_i$ and the creative's content vector $\mathbf{c}$ to $[0, 1]$. The content vector is a deterministic hashed representation of the creative title, biased for a small set of gender-, age-, and city-related keywords and then normalized; neither title in the case study contains these keywords. The activity level a_{ip} is the individual's use frequency on platform p. The targeting weight g_i^{target} equals the targeting strength 2.0 when the individual satisfies any condition listed in the audience setting (age band, gender, or city tier) and $1/2.0$ otherwise. Creator matching $g_i^{\text{creator}} = 1 + 0.5\,\text{clip}(\langle \mathbf{u}_i, \mathbf{k} \rangle, 0, 1)$ measures how close the individual's interests are to the creator's audience vector $\mathbf{k}$; when the creator's niche has a follower-profile prior, it is further multiplied by a gender-, age-, city-, and income-based weight normalized to mean one. The creator niche in the case study has no such prior, so this weight equals one. On RedNote, the platform audience prior π_{ip} multiplies 1.33 for women and 0.68 for men, 1.8 for tier-1 and tier-2 cities and 0.6 for tier-3 cities and below, and 1.4 for ages 15–44 and 0.6 for ages 55 and older. The exploration perturbation $\xi_i \sim \mathcal{U}(1 - 0.4\delta_p,\, 1 + 0.4\delta_p)$ is drawn with the scenario seed, with exploration strength $\delta_p = 0.65$ on RedNote.

The budget determines how many individuals are selected. The platform budget on the simulated sample is $b\rho_p / w_{\mathcal{P}}$, and the nominal number of impressions is $1000\, b\rho_p / (w_{\mathcal{P}}\, \text{CPM}_p)$, where the cost per thousand impressions on RedNote is CNY 48. OranSim selects individuals in decreasing order of e_i, up to the nominal number of impressions or the full sample size, whichever is smaller. For the CNY 40,000 budget of the case study, the sample budget is CNY 2,000 and the nominal number of impressions is 41,666.7, so 41,666 individuals are selected; scaled by $w_{\mathcal{P}}$, this corresponds to 833,333 nominal impressions and a represented reach of 833,320 people.

A.3 Individual Response and Paid Schedule

For each selected individual, the click and engagement probabilities are

$$\begin{aligned} p_i^{\text{click}} &= \sigma\big(\mathbf{w}_{\text{c}}^{\top} \mathbf{x}_i^{\text{c}} - 1.2 - \kappa_i + 0.7\,\epsilon_i\big)\,, \\ p_i^{\text{engage}} &= p_i^{\text{click}}\, \sigma\big(\mathbf{w}_{\text{e}}^{\top} \mathbf{x}_i^{\text{e}} - 1.5 + 0.5\,\epsilon_i\big)\,. \end{aligned} \tag{A.2}$$

The click features $\mathbf{x}_i^{\text{c}}$ are content matching, platform activity, targeting weight, creator matching, a fatigue proxy, openness, and a celebrity indicator, with weights $\mathbf{w}_{\text{c}} = (1.8, 1.2, 0.9, 0.4, -0.7, 0.25, 0.3)$; the fatigue proxy is the absolute value of the first response-state coordinate, clipped to $[0, 1]$. The engagement features $\mathbf{x}_i^{\text{e}}$ are content matching, platform

activity, targeting weight, creator matching, neuroticism, and openness, with weights $\mathbf{w}_{\mathrm{e}} = (1.4, 0.9, 0.6, 0.5, -0.3, 0.4)$. The term κ_i collects click penalties from compliance risk and AIGC labeling of the creative; both are zero for the two creatives in the case study. The latent response perturbation $\epsilon_i \sim \mathcal{N}(0,1)$ is drawn for every sampled individual with seed $\omega + 71{,}000$, so an individual uses the same perturbation in every scenario under a given seed. The mean engagement probability reported in the main text is the mean of p_i^{engage} over selected individuals.

The initial response mass entering propagation is $p_i^{\mathrm{click}} + 0.5\, p_i^{\mathrm{engage}}$. Individual masses are summed within segments and then distributed over days by the paid schedule: the share of day d is proportional to $\exp[-0.4(d-1)]$, normalized over 14 days, and divided equally among four six-hour steps. Under this schedule, day 1 receives 33.1% of the paid response mass and the first two days receive 55.3%.

A.4 Segment-Level Propagation

In Equation (2), the decay rate is $\beta = 0.9$ per day, the branching parameter is $r = 0.35$, and the step size is $\Delta t = 0.25$ days. The intergroup influence matrix $A \in \mathbb{R}^{60\times 60}$ is built from segment attributes. The unnormalized weight is 3.0 within a segment and, between segments $i \neq j$,

$$\tilde{A}_{ij} = 0.15 \exp(-0.5\, d_{ij}), \qquad d_{ij} = |\Delta_{\mathrm{age}}| + 0.8\, \mathbb{1}[\text{genders differ}] + 0.6\, |\Delta_{\mathrm{city}}|, \tag{A.3}$$

where Δ_{age} and Δ_{city} are differences in age-band and city-tier indices. Each row is then divided by its sum, so the influence weights of every source segment sum to one. The construction makes interpersonal influence stronger between similar people: influence is largest within a segment and weakens as age, gender, and city tier differ more. Studies of interpersonal communication find that people receive information and opinions more readily from others like themselves (Katz & Lazarsfeld, 1955), and agent-based marketing models represent such influence through local connections between consumers (Rand & Rust, 2011).

Because every row of A sums to one, summing Equation (2) over segments gives a recursion for the total excitation $z_\tau = \mathbf{1}^\top \mathbf{S}_\tau$:

$$z_{\tau+1} = \big(e^{-\beta \Delta t} + r\Delta t\big) z_\tau + \mathbf{1}^\top \boldsymbol{\mu}_\tau. \tag{A.4}$$

With the values above, $e^{-\beta\Delta t} + r\Delta t = 0.886 < 1$, so total excitation decays geometrically after the paid input ends. The total organic response is therefore determined by the paid input and the parameters r, β, and Δt, and A determines how it is distributed across segments.

A.5 Rollout Procedure

For a scenario s and seed ω, the rollout operator R_ϕ performs the following steps.

1. **Shared population.** Use the population of Appendix A.1; the scaling weight $w_{\mathcal{P}}$ and the scenario's budget and platform shares give the sample budget.
2. **Exposure.** Draw the exploration perturbation with seed ω, compute exposure scores by Equation (A.1), and select individuals.
3. **Individual response.** Read the latent perturbations ϵ_i under seed $\omega + 71{,}000$ and compute click probabilities, engagement probabilities, and initial response masses by Equation (A.2).
4. **Paid injection.** Sum initial response masses within segments and distribute them over 56 steps by the paid schedule, giving $\boldsymbol{\mu}_\tau$.
5. **Propagation.** Starting from $\mathbf{S}_0 = \mathbf{0}$, compute the organic response rate $\boldsymbol{\lambda}_\tau$ and the state $\mathbf{S}_{\tau+1}$ step by step by Equation (2).
6. **Aggregation.** Merge the four steps of each day into daily paid and organic response masses, scale them by $w_{\mathcal{P}}$, and accumulate them into $M_{14}(s,\omega)$ by Equation (4); record the selected individuals, mean content matching, mean engagement probability, and daily trajectories.

Scenarios differ only in the actions used in steps 2–6; the population, mechanism parameters, and per-seed latent perturbations stay fixed.

Table A.1: Simulator parameters.

Symbol	Value	Role	Section
N, $w_{\mathcal{P}}$	100,000; 20	Sample size and population scaling weight	A.1
–	2027	Population seed	A.1
–	$6 \times 2 \times 5 = 60$	Population segments	A.1
–	2.0; $1/2.0$	Targeting strength (matched; unmatched)	A.2
–	$1 + 0.5\,\mathrm{clip}(\cdot)$	Creator matching	A.2
π_{ip}	1.33/0.68; 1.8/0.6; 1.4/0.6	RedNote audience prior (gender; city tier; age)	A.2
δ_p	0.65	RedNote exploration strength	A.2
CPM_p	CNY 48	RedNote cost per thousand impressions	A.2
$\mathbf{w}_{\mathrm{c}}$; intercept	$(1.8, 1.2, 0.9, 0.4, -0.7, 0.25, 0.3)$; -1.2	Click model	A.3
$\mathbf{w}_{\mathrm{e}}$; intercept	$(1.4, 0.9, 0.6, 0.5, -0.3, 0.4)$; -1.5	Engagement model	A.3
–	0.7; 0.5	Latent perturbation scale (click; engagement)	A.3
–	0.5	Engagement weight in initial response mass	A.3
–	0.4	Daily decay of the paid schedule	A.3
β	0.9 day^{-1}	Excitation decay rate	A.4
r	0.35	Branching parameter	A.4
Δt	0.25 day	Step size	A.4
–	3.0; 0.15; (1, 0.8, 0.6)	Influence matrix: within-segment weight; cross-segment scale; age, gender, city distance coefficients	A.4

B Case Study: Controlled Skincare Campaign

This appendix records the controlled case of Section 5.2 in full. It follows the order of the main text: the frozen campaign specification, the reach, response tendencies, and cumulative response of all five options, their daily trajectories, the application of the collect-first rule, seed variation with reproducibility checks, mechanism controls, and parameter sensitivity.

B.1 Campaign Specification

Table B.1 lists the settings shared by all options. The audience setting lists conditions on age band, gender, and city tier, and targeting weights follow Appendix A.2. Both creatives use synthetic Chinese-language copy written by the authors; the table gives English translations. Table B.2 lists the five options: each single-factor option changes one field of the baseline, and the combined option changes the creative and the budget.

B.2 Reach, Response Tendencies, and Cumulative Response

Table B.3 lists the process quantities of the five options in the order of the rollout. Nominal impressions determine how many individuals are selected; the content matching and engagement probability of the selected individuals determine the initial response mass; the initial responses are injected according to the paid schedule, excite organic responses, and accumulate into the 14-day cumulative simulated response mass.

Options with the same budget have the same nominal impressions and reach, and their process quantities differ through the composition of the selected individuals. At CNY 40,000, the creator change moves the mean content match from 0.5009 to 0.5011 and the mean engagement probability from 0.7020 to 0.7015, with a cumulative simulated response mass of 1,833,903. Doubling the budget increases the number of selected individuals from 41,666 to 83,333. Under creative B, raising the budget from CNY 40,000 to CNY 80,000 moves the mean content match from 0.4971 to 0.4987, lowers the mean engagement probability from 0.7009 to 0.6799, and raises the cumulative simulated response mass from 1,832,978 to 3,604,845.

Table B.1: Campaign settings shared by all options.

Field	Setting
Platform and niche	RedNote; skincare
Publication time	2026-03-16 10:00
Outcome evaluation age γ	14 days
Audience setting	Age bands 15–24, 25–34, and 35–44; women; tier-1 to tier-3 cities; targeting strength 2.0
Platform shares $\boldsymbol{\rho}$	RedNote 100%
Schedule and propagation	Appendices A.3–A.4
Population and seeds	Shared population of Appendix A.1; seeds $\{0, \ldots, 29\}$
Creative A	Title: "Spring repair for sensitive skin: a triple-ceramide formula and a 28-day usage guide." Body: introduces the formula, suitable skin types, and morning and evening steps, and points to the full ingredient list. Topics: sensitive skin, ceramides, spring skincare. 30-second video; minimal visuals; calm music.
Creative B	Title: "Redness when the season changes? A 30-second look at a triple-ceramide repair routine." Body: invites viewers to take a skin-type test and claim a trial-size sample, and points to a staged 28-day routine. Topics: seasonal redness, sensitive skin, trial size. 30-second video; bright visuals; calm music.
Mid-tier creator	300,000 followers; interaction rate 4.5%
High-tier creator	1,200,000 followers; interaction rate 3.8%
Campaign objective	Budget cap CNY 80,000; compare predicted collects first and $\overline{M}_{14}$ for ties

Table B.2: Campaign options and the fields they change relative to the baseline.

Option	Creative	Budget (CNY)	Creator	Changed fields
s_0, baseline	A	40,000	mid-tier	–
s_c, creative change	B	40,000	mid-tier	creative
s_b, budget increase	A	80,000	mid-tier	budget
s_k, creator change	A	40,000	high-tier	creator
s_{cb}, creative and budget	B	80,000	mid-tier	creative, budget

B.3 Daily Paid and Organic Response

Figure B.1 shows the daily paid and organic responses of the five options. Paid response is highest on day 1 under the paid schedule and declines thereafter; organic response accumulates from the excitation of earlier responses, peaks on day 3, and then decays. Options with the same budget nearly coincide at the plotted scale, and on every day the paid and organic responses of the two CNY 80,000 options are about 1.96 times those of the CNY 40,000 options. For s_0, the paid and organic responses are 345,376 and 41,984 on day 1, 155,188 and 129,033 on day 3, and 1,905 and 6,041 on day 14.

B.4 Applying the Collect-First Rule

Equation (7) is applied in three steps. The budget cap of CNY 80,000 keeps all five options. Ranking by predicted collects gives 1,480 for s_k, 1,371 for s_0 and s_b, and 1,020 for s_c and s_{cb}; the budget is not a predictor input, so options with the same creative and creator but different budgets have identical predictions. Within each tied pair, the rule compares $\overline{M}_{14}$: s_b precedes s_0 (3,599,734 versus 1,834,405), and s_{cb} precedes s_c (3,604,845 versus 1,832,978). The complete order is s_k, s_b, s_0, s_{cb}, s_c, and the rule selects s_k. Table B.4 lists the change of each prediction relative to the baseline: the creator change raises all four predictions, and the creative change raises reads and comments while lowering likes and collects.

B.5 Seed Variation and Reproducibility Checks

Across the 30 shared seeds, the standard deviation of M_{14} lies between 992 and 1,299 for the five options. Differences between options are computed within each seed, where the population and latent response noise are shared. Table B.5 reports the mean paired differences with percentile intervals from 10,000 bootstrap resamples of the 30 seeds, drawn jointly for all options. The creative effect

Table B.3: Rollout quantities of the five options. Means ± standard deviations are over 30 shared seeds; reach is identical across seeds. Response masses are scaled by $w_{\mathcal{P}}$. Match and engagement are means over selected individuals.

(a) Exposure and individual response

Option	Nominal impressions	Sample reach	Represented reach	Content match	Engagement probability
s_0	833,333	41,666	833,320	0.5009 ± 0.0001	0.7020 ± 0.0007
s_c	833,333	41,666	833,320	0.4971 ± 0.0000	0.7009 ± 0.0007
s_b	1,666,667	83,333	1,666,660	0.4942 ± 0.0000	0.6783 ± 0.0004
s_k	833,333	41,666	833,320	0.5011 ± 0.0001	0.7015 ± 0.0007
s_{cb}	1,666,667	83,333	1,666,660	0.4987 ± 0.0000	0.6799 ± 0.0004

(b) Cumulative response over 14 days

Option	14-day paid response	14-day organic response	$\overline{M}_{14}$
s_0	$1{,}043{,}738 \pm 565$	$790{,}668 \pm 428$	$1{,}834{,}405 \pm 992$
s_c	$1{,}042{,}926 \pm 589$	$790{,}053 \pm 446$	$1{,}832{,}978 \pm 1{,}034$
s_b	$2{,}048{,}172 \pm 738$	$1{,}551{,}562 \pm 559$	$3{,}599{,}734 \pm 1{,}296$
s_k	$1{,}043{,}452 \pm 585$	$790{,}451 \pm 443$	$1{,}833{,}903 \pm 1{,}029$
s_{cb}	$2{,}051{,}080 \pm 739$	$1{,}553{,}765 \pm 560$	$3{,}604{,}845 \pm 1{,}299$

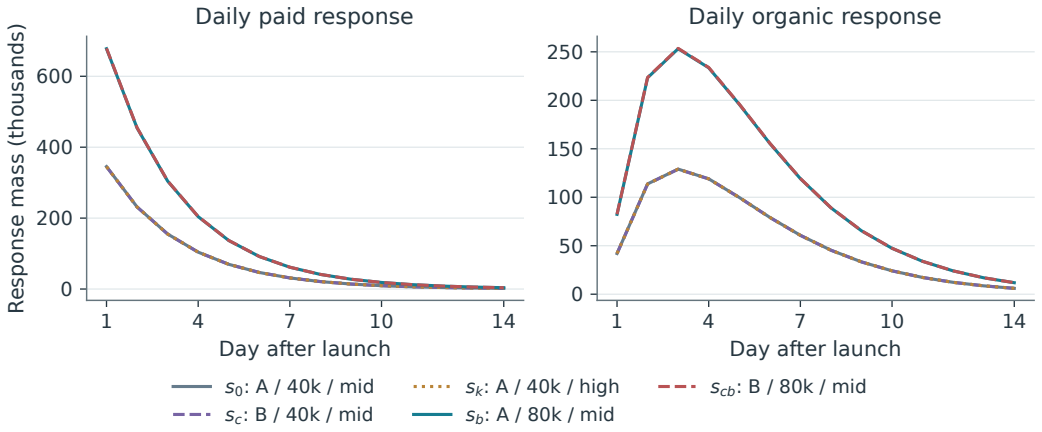

Figure B.1: Daily paid (left) and organic (right) response of the five options. Lines are means over 30 shared seeds and bands show ±1 standard deviation.

changes sign with the budget: creative B lowers M_{14} at CNY 40,000 and raises it at CNY 80,000. Relative differences divide by the comparator's mean, and per-CNY differences first divide each option's M_{14} by its budget. The intervals describe simulation randomness for the fixed population, mechanism parameters, and note predictions.

The frozen results include the following checks. Repeating the baseline with the same seed gives a maximum absolute difference of zero in both aggregate and daily outputs; an empty intervention that changes no field gives a zero difference; all five options share the same population hash; and all rollout summaries are finite. The scenario manifest records the fields each option changes relative to the baseline, matching Table B.2.

The predictions use the v3.1-pg predictor with title and body representations from its training encoder, text-embedding-3-small (first 768 dimensions, renormalized). The four vectors are stored with the results, so the predictions are reproduced without API access. The principal component transformation in the predictor was serialized with scikit-learn 1.8.0 and emits a version warning when loaded in the 1.6.1 runtime; inference completes, and the determinism and finite-output checks above pass.

Table B.4: Change in predicted engagement counts relative to the baseline s_0, with relative changes in parentheses.

Option	Reads	Likes	Collects	Comments
s_c, s_{cb}	+1,701 (+3.1%)	-193 (-8.7%)	-350 (-25.6%)	+28 (+53.2%)
s_k	+6,063 (+11.1%)	+463 (+21.0%)	+109 (+8.0%)	+27 (+52.0%)
s_b	0	0	0	0

Table B.5: Paired differences in M_{14} over the 30 shared seeds, with 95% bootstrap intervals.

Contrast	Mean difference	95% interval	Relative	Per CNY
$s_c - s_0$	$-1{,}427$	$[-1{,}510, -1{,}343]$	-0.08%	-0.08%
$s_k - s_0$	-502	$[-548, -456]$	-0.03%	-0.03%
$s_b - s_0$	1,765,329	[1,764,964, 1,765,680]	$+96.23\%$	-1.88%
$s_{cb} - s_b$	5,111	[5,056, 5,168]	$+0.14\%$	$+0.14\%$
$(s_{cb} - s_b) - (s_c - s_0)$	6,538	[6,436, 6,639]	$+0.36\%$	$+0.22\%$

The frozen results also record hashes of the campaign specification, model checkpoints, population, random-number configuration, and stored text representations.

B.6 Mechanism Controls

Two controls keep the number of selected individuals and each individual's latent response noise. Uniform selection replaces the exposure ranking by a fixed random order over individual identifiers and selects its first K_p individuals, so the sets selected under different budgets are nested; each selected individual keeps its own features. The mean-feature control keeps the individuals selected by the full rollout and replaces each response feature, after its transformations, by its mean over the selected audience before computing click and engagement probabilities. Both controls use the same paid schedule and propagation. Besides the five options, three options extend the comparison: s_l halves the baseline budget to CNY 20,000, and s_t and s_{tb} set every targeting weight to one at CNY 40,000 and CNY 80,000. The populations generated with seeds 2028 and 2029 repeat all comparisons.

Table B.6 reports the population of the main text. The 95% paired bootstrap intervals of the differences, from 10,000 resamples of the 30 seeds, lie within ± 0.04 percentage points of each estimate, and each percentage changes by less than 0.07 percentage points across the three populations. The full rollout exceeds uniform selection by 1.1%–8.5%, with the largest margin at the smallest budget, where the ranking is truncated earliest. Its difference from the mean-feature control stays between -0.21% and -0.14%. With neutral targeting, $\overline{M}_{14}$ at CNY 40,000 falls from 1,834,405 to 1,561,241, although all selected individuals still meet at least one audience condition (97.6% at CNY 80,000). The targeting weight enters both the exposure score and the click and engagement features (Appendix A.3), and setting it to one lowers the mean engagement probability from 0.7020 to 0.5368.

The propagation comparison fixes, for every seed, the 60 segment injections $\boldsymbol{\mu}_\tau$ of an option and replaces only A. The within-segment matrix is the identity. The uniform cross-segment matrix keeps each diagonal weight of A and spreads the remaining weight of its row equally over the other 59 segments. Table B.7 compares the daily distribution of organic response over segments with the distribution under A by the total variation distance $\frac{1}{2}\sum_g |p_g(d) - q_g(d)|$, averaged over the 14 days and 30 seeds. Total response differs from that under A by at most 2×10^{-7} in relative terms, the float32 rounding level, as Equation (A.4) implies. With within-segment propagation every segment has the same organic timing; under A, the time centroids of segment organic response range from 4.79 to 5.64 days. Option s_t and the other two populations change the distances and shares by less than 0.001 and the centroids by less than 0.02 days.

Table B.6: Full rollout versus the two equal-reach controls in the main-text population. Percentages give how much the full rollout's $\overline{M}_{14}$ exceeds each control, averaged over 30 seeds.

Option	Budget (CNY)	$\overline{M}_{14}$	vs. uniform selection	vs. mean features
s_0	40,000	1,834,405	$+6.61\%$	-0.17%
s_c	40,000	1,832,978	$+6.19\%$	-0.16%
s_b	80,000	3,599,734	$+4.59\%$	-0.19%
s_k	40,000	1,833,903	$+6.59\%$	-0.17%
s_{cb}	80,000	3,604,845	$+4.42\%$	-0.17%
s_l	20,000	933,383	$+8.49\%$	-0.14%
s_t, neutral targeting	40,000	1,561,241	$+4.45\%$	-0.20%
s_{tb}, neutral targeting	80,000	3,021,485	$+1.08\%$	-0.21%

Table B.7: Organic response of s_0 under three intergroup matrices with fixed initial segment responses.

Matrix	Mean daily TV distance from A	Share in targeted segments	Segment time centroids (days)
Structured A	0	96.05%	4.79–5.64
Within-segment	0.246	100%	5.02
Uniform cross-segment	0.067	94.26%	4.75–5.42

B.7 Parameter Sensitivity

Each setting changes one value of Table A.1 and keeps the others: the decay rate β, the branching parameter r, the targeting strength, the RedNote exploration strength δ_p, or both response weight vectors $\mathbf{w}_{\mathrm{c}}$ and $\mathbf{w}_{\mathrm{e}}$, scaled by a common factor. Table B.8 gives the budget-doubling ratios. The rankings of the five options are identical in all eleven settings: $s_k, s_b, s_0, s_{cb}, s_c$ by the collect-first rule; $s_{cb}, s_b, s_0, s_k, s_c$ by $\overline{M}_{14}$; and $s_0, s_k, s_c, s_{cb}, s_b$ by response per CNY. Changing β or r multiplies every option's total by the same factor, because all options share the paid schedule and the aggregate recursion (A.4) is linear, so the ratios stay unchanged. The collect-first ranking uses the fixed note predictions. The settings are a local stress test around the design values.

Table B.8: Ratio of $\overline{M}_{14}$ to the baseline s_0 for the two CNY 80,000 options under single-parameter changes.

Parameter	Values	s_b/s_0	s_{cb}/s_0
Design values	Table A.1	1.962	1.965
Decay rate β	0.6; 1.2	1.962; 1.962	1.965; 1.965
Branching parameter r	0.2; 0.5	1.962; 1.962	1.965; 1.965
Targeting strength	1.5; 2.5	1.950; 1.972	1.953; 1.974
Exploration strength δ_p	0.325; 0.975	1.963; 1.962	1.966; 1.965
Response weight scale	0.8; 1.2	1.953; 1.973	1.956; 1.975

C RedNote Engagement Prediction

C.1 Corpora and Features

The two corpora support different comparisons. The first contains 39,000 note snapshots and compares the configurations without note age (v3) and with note age (v3.1-pg) on the same five folds, fixed after shuffling the notes. The second contains 12,154 notes from 13 niches, each with at least 100 reads, and is used for the structured holdout experiments.

For the 39,000-note comparison, titles and bodies are embedded separately with OpenAI text-embedding-3-small. Each 1,536-dimensional output keeps its first 768 dimensions and is renormalized to unit length; the two vectors are concatenated into 1,536 dimensions and reduced to 128 principal components. The resulting predictor inputs contain 182 features for v3 and 183 for v3.1-pg. Each outcome regressor uses up to 500 LightGBM trees, learning rate 0.05, maximum depth 6, 32 leaves, a minimum of 20 examples per leaf, and L1 and L2 penalties of 0.3. Five-fold cross-validation uses shared shuffled folds with seed 42 and early stopping after 30 rounds; the regressor seed is 0.

Each structured holdout split refits preprocessing and predictors. Titles and bodies each receive a 768-dimensional deterministic hashed representation; the two are concatenated and reduced to at most 64 dimensions by randomized principal component analysis fitted on the training set only. The remaining features are creator follower count, content duration, media type, niche, the 30 most frequent topics in the training set, cyclic encodings of publication hour and weekday, and note age. Each outcome has its own LightGBM regressor trained on $\log(1+y)$, and $R^2_{\log}$ and $\mathrm{RMSE}_{\log}$ are computed on the log scale. The structured holdout regressors use 300 trees, learning rate 0.05, maximum depth 6, 32 leaves, a minimum of 20 examples per leaf, and L1 and L2 penalties of 0.3.

C.2 Holdout Splits

Table C.1 lists the three holdout splits. The temporal split trains on the earliest 85% of notes by publication time and tests on the latest 15%. The unseen-creator split groups notes by creator; none of the 1,642 creators in the test set appears in training. The leave-one-niche-out split holds out each niche in turn, and the metrics are computed on the pooled test predictions of the 13 folds. Table C.2 lists the number of notes in each niche.

C.3 Count-Scale Errors

Table C.3 reports the raw-count mean absolute error (MAE) and the ratio of the mean predicted count to the mean observed count. For the structured holdouts, both come from the stored MAE and ten-bin calibration aggregates, and mean predictions are 18%–43% of the observed means in all three splits and four outcomes. For the 39,000-note predictor used in the campaign case, repeating the five-fold protocol of Table 1 on a re-extraction of the corpus reproduces its $R^2_{\log}$ to within 0.003; out-of-fold mean predictions are 50%–59% of the observed means, and their medians are 98%–107% of the observed medians. The regressors fit $\log(1+y)$ and report $\exp(\hat{z})-1$, which lies below the conditional mean of a right-skewed count.

Table C.1: Structured holdout experiments on the 12,154-note corpus.

Split	Train / test	Metric	Reads	Likes	Collects	Comments
Latest-time	10,331 / 1,823	$R^2_{\log}$	0.253	0.361	0.329	-0.007
		$\mathrm{RMSE}_{\log}$	1.879	1.972	1.831	1.525
Unseen creator	10,310 / 1,844	$R^2_{\log}$	0.559	0.550	0.513	0.347
		$\mathrm{RMSE}_{\log}$	1.506	1.736	1.688	1.463
Leave-one-niche-out	12,154 / 13 folds	$R^2_{\log}$	0.503	0.513	0.456	0.333
		$\mathrm{RMSE}_{\log}$	1.625	1.811	1.781	1.500

Table C.2: Niche composition of the holdout corpus.

Niche	Notes	Niche	Notes	Niche	Notes
Fashion	1,466	Home	1,202	Automotive	819
Parenting	1,374	Pets	1,062	Fitness	618
Food	1,309	Electronics	987	Health care	437
Skincare	1,221	Education	970	Travel	410
Makeup	279				

C.4 Note-Age Reference

In training, note age counts the days from publication to the date the predictor is fitted, capped at 180 days, whereas the engagement counts were collected earlier. To measure how this reference affects temporal evaluation, a re-extracted sample of 12,154 notes with the size and niche composition of the holdout corpus is split by whole publication dates into 9,265 training notes (March 17–26, 2026), 1,749 calibration notes (March 27), and 1,140 test notes (March 28 to April 19). The recorded update dates of these rows range from March 24 to April 26. Hashed-text LightGBM regressors are refitted with note age counted either to the extraction date, September 25, 2026, which caps 96.8% of ages at 180 days, or to the date of the original holdout run, May 5, 2026, which caps none.

Changing only this reference lowers test MAE by 43%–58% across the four outcomes (Table C.4). Under both references the test $R^2_{\log}$ is negative; this split trains on ten publication days and tests on the following three weeks. A log-linear recalibration fitted on the calibration day raises $R^2_{\log}$ under the May reference to 0.158, 0.258, 0.227, and -0.068, while mean predictions remain 70%–85% below the observed means. A multiplicative recalibration that matches the calibration-day mean count increases MAE for all four outcomes.

Table C.3: Count-scale errors of the RedNote predictors. Ratio is the mean predicted count divided by the mean observed count.

	Reads		Likes		Collects		Comments	
Split	MAE	Ratio	MAE	Ratio	MAE	Ratio	MAE	Ratio
39,000-note five-fold CV	62,594	0.59	3,116	0.54	1,145	0.52	226.6	0.50
Latest-time	17,305	0.24	1,025	0.18	332	0.21	42.7	0.39
Unseen creator	43,423	0.43	2,232	0.36	876	0.31	145.0	0.37
Leave-one-niche-out	50,018	0.40	2,329	0.34	848	0.30	178.9	0.31

Table C.4: Temporal diagnostic on a re-extracted 12,154-note sample under two note-age references. MAE is on the count scale; the last two columns use the May reference with log-linear recalibration.

	Test MAE			Log-linear recalibration	
Outcome	September reference	May reference	$R^2_{\log}$, May	MAE	Mean bias
Reads	29,306	13,446	-0.508	11,332	-82.0%
Likes	1,182	676	-0.260	586	-84.3%
Collects	507	291	-0.215	264	-85.3%
Comments	103.3	43.8	-1.001	31.6	-70.1%

D Policy Value and Audience Ranking

D.1 Off-Policy Estimators and Identification

All four estimators of Section 4.2 use the importance weight $w_i = \pi(a_i \mid x_i)/\mu_i$. DR and Switch-DR (Dudík et al., 2011; Wang et al., 2017b) also use an outcome model $\hat{q}(a)$ that shrinks the mean reward of each action in the earlier part of a log toward that part's mean reward,

$$\hat{q}(a) = \frac{\sum_{j:a_j=a} y_j + \beta_q \bar{y}}{|\{j : a_j = a\}| + \beta_q}, \qquad \beta_q = 20, \tag{D.1}$$

where the sums run over the earlier part and $\bar{y}$ is its mean reward; on Open Bandit, $\hat{q}$ is indexed by action and recommendation position. The model value of the target policy in context x is $V_\pi(x) = \sum_a \pi(a \mid x)\hat{q}(a)$. With residual $\delta_i = y_i - \hat{q}(a_i)$,

$$\hat{V}_{\mathrm{DR}}(\pi) = \frac{1}{n}\sum_{i=1}^{n}\big[V_\pi(x_i) + w_i\delta_i\big], \qquad \hat{V}_{\text{Switch-DR}}(\pi) = \frac{1}{n}\sum_{i=1}^{n}\big[V_\pi(x_i) + w_i \mathbb{1}\{w_i \le \tau\}\,\delta_i\big] \tag{D.2}$$

Switch-DR applies the residual correction only to records whose weight does not exceed τ. The KuaiRand-Pure experiment uses $\tau = 100$ and the Open Bandit experiment $\tau = 5$, with the logged propensity scores and recommendation positions of the Open Bandit Dataset (Saito et al., 2020). Because the outcome model does not depend on user context, the model value is the same in every context.

Importance weighting identifies the policy value under three conditions: (1) overlap, $\pi_0(a \mid x) > 0$ for every context x and every action a with $\pi(a \mid x) > 0$; (2) unconfoundedness, potential outcomes are independent of the logged action given the context; and (3) the stable unit treatment value assumption, under which a unit's potential outcome depends only on its own action. Under these conditions,

$$V(\pi) = \mathbb{E}_{x,\,a\sim\pi}[Y(a)] = \mathbb{E}_x\Big[\sum_a \pi(a \mid x)\,\mathbb{E}[Y \mid x, a]\Big] = \mathbb{E}_{x,\,a\sim\pi_0}\Big[\frac{\pi(a \mid x)}{\pi_0(a \mid x)}\,Y\Big], \tag{D.3}$$

where the second equality uses unconfoundedness and the third uses overlap (Li et al., 2011). KuaiRand-Pure inserts videos sampled uniformly from its candidate pool, and the Open Bandit Dataset records the assignment probability of every logged record.

The target policy and outcome model come from an earlier part of each log, and the estimates use the later part. KuaiRand-Pure splits its random log by calendar date. April 22–29, 2022, with

Table D.1: Off-policy estimates on the later part of each log. ESS, mean importance weight, and weight range describe how each log supports the target policy; the mean weight has expectation one when the weights use the probabilities with which the log was generated.

Data / logging policy	IPS	SNIPS	DR	Switch-DR	ESS	Mean weight	Weight range
KuaiRand / uniform	0.0932	0.0879	0.0883	0.0828	1,454	1.060	$[0.10, 682.57]$
Open Bandit / random	0.00902	0.00998	0.00915	0.00280	685	0.904	$[0.10, 7.30]$
Open Bandit / BTS	0.00209	0.00216	0.00258	0.00396	611	0.969	$[0.0027, 109.94]$

337,933 records, selects the ten most-logged videos among the 7,583 candidates, breaking ties by identifier, and fits $\hat{q}$; April 30 to May 8 is evaluated after removing 153 records whose timestamps precede the last construction event, which leaves 847,973. Open Bandit splits each log at its median timestamp. The two earlier halves together select one set of ten actions, each earlier half fits its own $\hat{q}$, and each later half is evaluated after the last construction timestamp of both logs, leaving 5,000 random and 4,964 BTS records. A selected action has target probability $0.1/80 + 0.9/10$ at every recommendation position, and any other action $0.1/80$. Intervals are percentile intervals from 1,000 bootstrap resamples of users on KuaiRand and of equal-timestamp groups on Open Bandit, conditional on the fixed target policy and outcome model.

When the weights use the probabilities with which the log was generated, the mean importance weight has expectation one. It is 1.060 with interval $[1.005, 1.114]$ on KuaiRand, 0.904 with $[0.841, 0.969]$ on the random Open Bandit log, and 0.969 with $[0.903, 1.045]$ on the BTS log. The KuaiRand logging probability $1/7{,}583$ follows the uniform sampling of randomly inserted videos described with the dataset (Gao et al., 2022). The candidate pool of the random log changes over time. Treating each video as a candidate from the first to the last day on which it is logged, the pool shrinks from 7,475 videos on April 23 to 7,214 on May 8: 369 videos stop appearing before May 8, and 72 first appear after April 23. Per-video exposure counts over the whole log have a dispersion index of 6.08, against 1.00 under uniform sampling from the 7,583 videos; among the 7,198 videos logged by April 23 and still logged on May 8, the index is 1.17, and the correlation between their construction and evaluation counts is 0.05. If each day's videos are drawn uniformly from that day's pool, weights computed with $1/7{,}583$ have an expected mean of 1.036 for the target policy. The ten videos it selects are in the pool on every evaluation day, and the daily pools cover 99.61% of its probability on average over the evaluated records. Replacing $1/7{,}583$ by the reciprocal of each day's pool size lowers the observed mean weight to 1.019 and gives SNIPS 0.0880, compared with 0.0879 under $1/7{,}583$. Building the KuaiRand target policy from the standard log of April 8–21 instead, and evaluating it on the whole random log, gives SNIPS 0.2545 with interval $[0.2303, 0.2762]$, an effective sample size of 1,596, and a mean weight of 0.804 with interval $[0.767, 0.845]$. This target policy differs from the one in Table D.1, so the two values describe different policies.

D.2 Uplift Metrics

Test customers are sorted by predicted effect $\hat{\tau}(x_i)$ in decreasing order. Let T_m be the first m customers, with treatment and control members $T_m^{(1)}$ and $T_m^{(0)}$. The gain curve is

$$U(m) = \left[\bar{y}(T_m^{(1)}) - \bar{y}(T_m^{(0)})\right] m. \tag{D.4}$$

Let $A(\hat{\tau})$ be the area under the gain curve, A_{rand} the area under the line from $(0, 0)$ to $(n, U(n))$, and A_{perfect} the area under the curve of the best ordering constructed from observed outcomes. The normalized AUUC is $(A(\hat{\tau}) - A_{\mathrm{rand}})/(A_{\mathrm{perfect}} - A_{\mathrm{rand}})$, which is 0 for a random ordering and 1 for the best ordering. The Qini metric replaces $U(m)$ with

$$Q(m) = \sum_{i \in T_m} y_i \mathbb{1}\{a_i = 1\} - \Big(\sum_{i \in T_m} y_i \mathbb{1}\{a_i = 0\}\Big) \frac{|T_m^{(1)}|}{|T_m^{(0)}|} \tag{D.5}$$

and is normalized in the same way. Uplift at a reach fraction is defined in Section 4.3. The three metrics match `uplift_auc_score`, `qini_auc_score`, and `uplift_at_k` with `strategy="overall"` in scikit-uplift (Surry & Radcliffe, 2011; Renaudin & Martin, 2021).

Table D.2: Audience ranking on X5 RetailHero. Entries are means ± standard deviations over three seeds; uplift columns are headed by the reach fraction. The highest mean in each column is in bold.

(a) Aggregate ranking metrics

Method	AUUC	Qini
Random	0.0040 ± 0.0016	0.0025 ± 0.0009
Two-model logistic	$\mathbf{0.0217 \pm 0.0019}$	0.0140 ± 0.0013
S-Learner	0.0125 ± 0.0038	0.0082 ± 0.0025
T-Learner	0.0155 ± 0.0054	0.0105 ± 0.0036
X-Learner	0.0213 ± 0.0030	$\mathbf{0.0145 \pm 0.0021}$
DRLearner	0.0202 ± 0.0005	0.0141 ± 0.0004
CausalForestDML	0.0112 ± 0.0048	0.0074 ± 0.0033

(b) Uplift at 5%, 10%, and 20% reach

Method	Uplift@5%	Uplift@10%	Uplift@20%
Random	0.0397 ± 0.0133	0.0385 ± 0.0100	0.0379 ± 0.0038
Two-model logistic	0.0861 ± 0.0073	0.0791 ± 0.0066	0.0664 ± 0.0041
S-Learner	0.0573 ± 0.0320	0.0676 ± 0.0144	0.0639 ± 0.0110
T-Learner	0.1337 ± 0.0069	0.0858 ± 0.0182	0.0713 ± 0.0053
X-Learner	0.1507 ± 0.0182	0.1045 ± 0.0118	$\mathbf{0.0875 \pm 0.0030}$
DRLearner	$\mathbf{0.1933 \pm 0.0287}$	$\mathbf{0.1306 \pm 0.0160}$	0.0773 ± 0.0064
CausalForestDML	0.0890 ± 0.0139	0.0737 ± 0.0091	0.0576 ± 0.0042

(c) Uplift at 30% and 50% reach

Method	Uplift@30%	Uplift@50%
Random	0.0388 ± 0.0025	0.0374 ± 0.0019
Two-model logistic	0.0585 ± 0.0045	0.0492 ± 0.0043
S-Learner	0.0549 ± 0.0037	0.0448 ± 0.0029
T-Learner	0.0606 ± 0.0038	0.0465 ± 0.0050
X-Learner	$\mathbf{0.0668 \pm 0.0046}$	$\mathbf{0.0498 \pm 0.0002}$
DRLearner	0.0639 ± 0.0026	0.0462 ± 0.0034
CausalForestDML	0.0470 ± 0.0004	0.0409 ± 0.0038

D.3 Case Study: Audience Ranking under Limited Reach

In X5 RetailHero, the advertiser sends a marketing message to randomly selected customers, with a treatment fraction of 0.4998. A campaign often reaches only part of its customers, so this case compares which ranking method yields the largest incremental response at each reach fraction. For each seed in $\{42, 137, 256\}$, the data are split 75%/25% into training and test sets, stratified jointly by treatment and outcome; all seven methods share the training and test customers, and preprocessing is fitted on the training set only. S-Learner, T-Learner, X-Learner, and DRLearner use 80-tree forests as base learners, CausalForestDML uses 200 trees, and two-model logistic regression fits separate logistic regressions to the treatment and control groups. All forests use a minimum leaf size of 50. The two logistic regressions and the logistic propensity model use the lbfgs solver with at most 500 iterations. DRLearner uses three-fold cross-fitting with separate 80-tree regression and final forests. CausalForestDML uses 80-tree random-forest classifiers for its outcome and treatment nuisance models, three-fold cross-fitting, and 200 causal trees. Table D.2 reports all methods, and Figure D.1 converts uplift into estimated incremental responses per 10,000 customers reached. DRLearner is highest when the campaign reaches the top 5% or 10% of customers, and X-Learner is highest from 20% to 50%; as the reach fraction grows, the incremental responses of all methods move closer to random ranking.

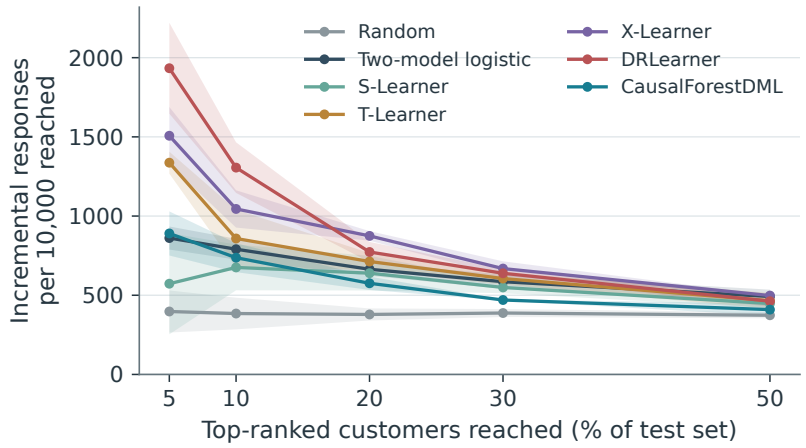

Figure D.1: Estimated incremental responses per 10,000 customers reached on X5 RetailHero, by the fraction of top-ranked customers reached. Lines are means over three seeds and bands show ± 1 standard deviation.

E Paired Counterfactual Scoring

E.1 Scenario Generator and Splits

The scenarios come from the synthetic scenario generator in the supplementary material, run with seed 2027 to produce 20,000 scenarios. The treatment arm of each scenario is determined jointly by its budget tier and creator tier, arm $= \min\big(3, \lfloor \text{budget tier}/2 \rfloor + 2 \cdot \mathbb{1}[\text{creator tier is mid or above}]\big)$. The alternative setting moves the scenario to arm $(\text{arm} + 1) \bmod 4$ and multiplies the budget by that arm's factor (0.6, 1.0, 1.4, and 1.8 for arms 0–3), keeping the other creator attributes fixed. The generator computes exposure, clicks, conversions, and revenue under both settings. In lexical order of scenario identifiers, the first 16,000 scenarios form the training set and the next two blocks of 2,000 form the validation and test sets, with no overlap. An identity intervention that keeps the arm and budget unchanged reuses the same computed potential outcome, so its effect is exactly zero.

The feature-substitution baseline uses platform, niche, budget tier, creator tier, treatment arm, log budget, log follower count, and creator engagement rate. Each outcome has its own LightGBM regressor (300 trees, learning rate 0.05, maximum depth 6, 32 leaves, at least 20 samples per leaf, and L1 and L2 regularization of 0.3) trained on $\log(1 + y)$.

E.2 Full Metrics and Segment-Level Effects

Table E.1 extends Table 4 with the mean absolute errors of factual outcomes, counterfactual outcomes, and individual effects, and with the rank correlation of individual effects. The grouped CATE R^2 is computed over 18 groups: four budget tiers, four creator tiers, and ten niches. Figure E.1 plots the generated and predicted effects of these 18 groups.

Table E.1: All metrics of the LightGBM feature-substitution baseline on the 2,000 paired synthetic test scenarios. Counterfactual normalized MAE divides the counterfactual MAE by the mean absolute alternative outcome.

(a) Factual and counterfactual outcome prediction

Outcome	Factual R^2	Factual MAE	Counterfactual MAE	Counterfactual normalized MAE
Exposure	0.964	18.24	61.14	0.219
Clicks	0.892	0.477	0.810	0.256
Conversions	0.762	0.0055	0.0121	0.582
Revenue	0.667	1.729	2.657	0.439

(b) Individual and grouped effect estimation

Outcome	Individual-effect MAE	Individual-effect sign accuracy	Individual-effect rank correlation	Grouped CATE R^2
Exposure	63.46	0.955	0.544	0.457
Clicks	0.920	0.617	0.429	0.355
Conversions	0.0138	0.930	0.450	0.531
Revenue	3.230	0.552	0.233	0.003

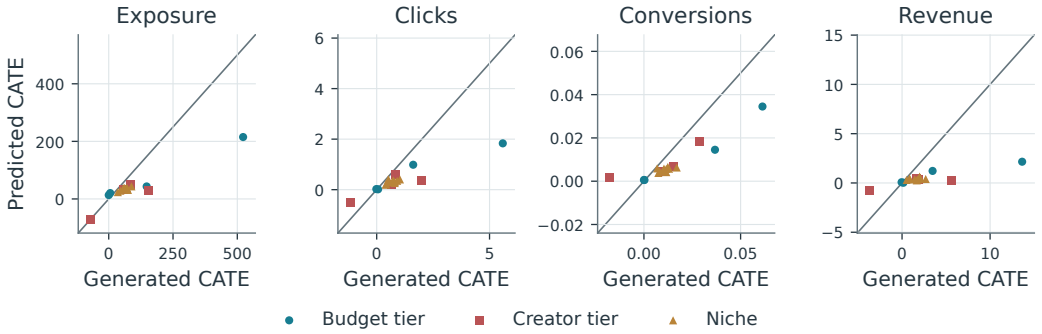

Figure E.1: Grouped CATE of the LightGBM feature-substitution baseline. Each point is a budget tier, creator tier, or niche; the horizontal axis is the within-group mean of generated effects and the vertical axis the within-group mean of predicted effects. The diagonal marks equality.

F Reproducibility Materials

F.1 Deterministic Output Snapshot

The output snapshot contains 3,000 scenario scoring records, 1,000 for each of seeds 42, 137, and 256, generated by the v3.1-pg predictor. The scenarios cover 15 niches, three creator tiers (nano, mid, and head), three budget tiers, two exposure regimes (platform recommendation and random), and image and video posts. Each record contains the scenario fields (niche, title, post type, duration, image count, creator tier, follower count, budget tier and amount, exposure regime, and note age), the four predicted engagement counts and their per-follower rates, and the predictor version and seed. Records also carry an exposure field set to eight times the predicted reads, which this paper does not use. Text representations are deterministic hashes; regenerating the snapshot with the same seeds and record count produces byte-identical files. Table F.1 summarizes the four predictions.

F.2 Supplementary and Private Materials

The supplementary material contains the simulator code, the synthetic scenario generator and its scenarios, and the scripts and aggregate results of the KuaiRand-Pure, Open Bandit, X5, and paired counterfactual experiments. The note predictors are trained on a private RedNote corpus, and Appendices B and C report the aggregate results of the RedNote experiments.

Table F.1: Predicted engagement counts in the output snapshot.

Outcome	Mean	Standard deviation
Reads	119,938.78	51,215.30
Likes	4,605.42	2,537.27
Collects	768.90	374.68
Comments	355.15	205.73