Proceedings of ICFTBA 2026 Symposium: Strategic Management & Business Analytics: Data-Driven Performance Models
DOI: 10.54254/2754-1169/2026.LD38497

Platform Attention Allocation under Adolescent Welfare Constraints: A Multi-Objective Mixed-Integer Programming Approach

Yuxuan Ren

The Second High School Attached to Beijing Normal University, Beijing, China
bam_gamma@outlook.com

Abstract. Video platforms allocate recommendation exposure primarily to maximize engagement, potentially creating behavioural welfare risks for adolescent users. Existing safeguards, such as screen-time reminders and forced interruptions, are largely ex post and do not alter the underlying allocation logic. This study formulates adolescent video recommendation as a constrained allocation problem and develops a multi-objective mixed-integer linear programming (MOMILP) framework that jointly incorporates predicted commercial utility, modeled behavioural risk, educational exposure, and content diversity. Using controlled synthetic experiments, this study compare the proposed framework with a myopic commercial baseline and a post-hoc hard-filtering baseline. Under the reported parameter setting, the proposed model reduces aggregate modeled risk by 58.9% relative to the commercial baseline while retaining 68.5% of its predicted commercial value at a nearly identical delivery cost. Compared with post-hoc filtering, joint optimization retains substantially more commercial value while coordinating multiple governance objectives within the allocation process. Sensitivity, solver-gap, and diversity analyses further demonstrate the trade-offs among commercial utility, modeled risk, computational performance, and exposure structure. These findings illustrate how adolescent-protection objectives can be embedded directly into recommendation allocation rather than added only as post-exposure controls. Results are based on synthetic data and should not be interpreted as causal estimates of real-world welfare improvement.

Keywords: platform attention allocation, recommendation exposure, adolescent welfare, commercial-utility-modeled-risk trade-off, multi-objective integer programming

1. Introduction

Video platforms use engagement proxies, such as clicks, watch time, retention, not only to predict user preferences but also to allocate scarce attention. However, these short-term commercial objectives incentivize platforms to concentrate exposure on high-engagement content. "Adolescents may be particularly vulnerable because their self-regulation capacities are still developing. For example, Xu et al. [1] find that daily short-video use predicts declines in academic delay of gratification and academic achievement in both adolescent cohorts. Against this background,

© 2026 The Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

DOI: 10.54254/2754-1169/2026.LD38497

negative effects are not necessarily internalized in the platform's immediate objectives. This incentive misalignment creates a fundamental conflict between platforms' commercial utility maximization and adolescent welfare.

Regulatory constraints aim to mitigate this conflict; platforms usually achieve these requirements by pop-up reminders and usage limits. These safeguards mainly operate after exposure has been generated. While easy to implement, such ex-post governance mechanisms do not fundamentally alter the allocation logic of exposure. The limitation lies in an objective that prioritizes predicted immediate commercial utility without coordinating governance constrains. This leads to the following research question: How can a platform integrate welfare-oriented constraints into its recommendation decisions to balance commercial utility against behavioural risk, while ensuring educational value and content diversity?

To address this gap, this paper formulates recommendation as a constrained allocation problem and develops a multi-objective mixed-integer linear programming (MOMILP) framework. The model explicitly integrates behavioural risk control, educational exposure, and content diversity into the decision layer. By embedding these dimensions as constraints and objectives, the framework shifts welfare considerations from ex-post correction to ex-ante allocation, enabling coordinated and globally feasible recommendation outcomes.

Accordingly, this paper makes three contributions. First, it reframes adolescent video recommendation as a platform constrained attention-allocation problem rather than a purely predictive ranking task. Second, it develops a MOMILP framework that jointly represents predicted commercial utility, modeled behavioural risk proxy, educational exposure, and content diversity under explicit governance constraints. Third, through controlled synthetic experiments, it quantifies the trade-off between commercial utility and modeled risk and examines its sensitivity, robustness, and computational feasibility.

Notation convention. Uppercase italic letters denote sets (e.g., I, J, and K), governance thresholds, or capacity parameters rather than random variables, since the present model is deterministic. Lowercase italic letters denote indices, parameters, and scalar decision variables. Subscripts i, j, and k index users, videos, and content categories, respectively. Superscripts + and – on slack variables denote exceedance and shortfall. Symbol definitions and units are reported in Appendix Table A1; generation rules for selected synthetic parameters are reported in Appendix A.2.

2. Literature review

2.1. Adolescent digital consumption and potential welfare costs

A growing body of research examines the relationship between digital technology use and adolescent welfare, identifying several potential welfare concerns associated with intensive or problematic consumption. These concerns include psychological distress, depressive symptoms, and anxiety, as well as patterns of compulsive or poorly self-regulated use [2, 3]. Importantly, these outcomes should not be interpreted as interchangeable measures of welfare. The former primarily represent psychological or mental-health outcomes, whereas the latter reflects a behavioral pattern through which excessive digital consumption may generate long-term welfare risks. At the same time, existing evidence suggests that average effects are often small in magnitude and highly heterogeneous across individuals and contexts. Large-scale empirical studies find only weak negative associations between digital use and welfare, cautioning against overstating aggregate harms [4]. Similarly, the causal impact of digital exposure remains difficult to identify, with

DOI: 10.54254/2754-1169/2026.LD38497

outcomes varying substantially depending on individual characteristics, usage patterns, and content types [5, 6].

Despite these mixed findings, several behavioral mechanisms provide a theoretical basis for welfare-relevant externalities. In particular, habit formation, self-control problems, and biased beliefs about future consumption can lead individuals to consume more digital content than is optimal from their own long-term perspective [3]. For adolescents, whose cognitive control and intertemporal decision-making capacities are still developing, such mechanisms may be especially pronounced. Excessive viewing, sleep displacement, and attention fragmentation can therefore arise even when short-term engagement appears privately optimal. These dynamics imply that observed usage behavior may not coincide with welfare-maximizing choices. In this study, this study focuses on poorly self-regulated and potentially addictive patterns of consumption. The model uses a synthetic addiction-risk proxy to represent the welfare-related risk associated with excessive digital consumption. The proxy should therefore be interpreted as a modeled behavioral risk indicator, rather than a direct measure of addiction or welfare.

Taken together, the literature suggests that digital consumption may generate heterogeneous welfare risks through behavioral frictions such as self-control problems. This study therefore model poorly self-regulated consumption as a potential behavioral internality and use a synthetic risk proxy to examine how such concerns can be incorporated into platform allocation decisions.

2.2. Platform engagement incentives and welfare trade-offs

Digital platforms operate under strong incentives to maximize user engagement, commonly measured by metrics such as clicks, watch time, and retention. While these metrics serve as practical proxies for user interest, they do not necessarily align with user welfare. A growing theoretical literature highlights that when platforms control the supply of information and optimize for engagement, they may induce excessive consumption and reduce overall welfare. In particular, Hébert and Zhong [7] show that maximizing user attention can lead to overexposure to highly engaging content, even when such exposure is not welfare-improving.

Empirical and theoretical evidence also suggests that observed engagement need not coincide with welfare-maximizing choices. Facebook deactivation experiments document a divergence between continued platform use and experienced well-being [8], while models of digital addiction emphasize that usage may reflect self-control failures [3]. More broadly, recommender systems can improve matching efficiency but may also steer users toward highly consumable content, particularly when observed engagement reflects behavioral mistakes [9, 10].

Taken together, this literature suggests that engagement-maximizing platforms may allocate attention in ways that deviate from welfare-improving outcomes. The evidence identifies a possible misalignment between commercial engagement objectives and user welfare; it does not identify a welfare optimum for adolescent short-video recommendation. This gap motivates a constrained-allocation model in which engagement-related value is optimized subject to explicit welfare-relevant constraints, rather than treating observed usage as a sufficient statistic for welfare.

2.3. Responsible recommendation and internal governance

To address the limitations of purely engagement-driven recommendation, recent research has increasingly explored responsible recommendation objectives, including fairness, diversity, popularity-bias mitigation, and user well-being. Milano, Taddeo, and Floridi [11] show that these objectives correspond to a broader set of ethical challenges, and that user-centred approaches often

DOI: 10.54254/2754-1169/2026.LD38497

leave the interests of other stakeholders—including content providers and society at large—under-specified. These objectives are often operationalized through distinct metrics or design interventions. Stray et al. [12] synthesize how human values are translated into metrics, ranking and product design, long-term objectives, and evaluation protocols. Complementing this operationalization, Helberger, Karppinen, and D'Acunto [13] treat exposure diversity as a design principle: recommenders can be engineered to broaden information exposure rather than to maximize engagement alone.

For short-video platforms, however, stating these objectives is not sufficient: ethical principles must become design choices, and those choices must become allocation rules that determine which users are matched to which content under limited attention. Constrained matching provides one route for this step. Mladenov et al. [14] show that long-term social welfare can be placed upstream of recommendation by formulating allocation as an optimal matching problem subject to explicit constraints, rather than relying on myopic engagement maximization. This connection between welfare and allocation supplies a template for internal governance: well-being-related limits can enter the matching stage itself, instead of remaining post-hoc evaluation criteria.

2.4. Constrained allocation and multi-object optimization

Operations research provides a natural framework for translating responsible-recommendation objectives into explicit allocation rules. Multi-objective recommender systems commonly involve competing quality metrics and stakeholder interests [15, 16], while Pareto-based methods, weighted-sum scalarization, and the ε-constraint method provide systematic tools for examining such trade-offs [17, 18]. These approaches allow assignment decisions, capacity limits, and governance requirements to be represented explicitly rather than embedded implicitly in a ranking score.

What remains underdeveloped is a unified and auditable allocation model that jointly represents the four objectives—commercial utility, behavioural-risk control, educational exposure, and content diversity—under explicit operational constraints. The present problem is well suited to a multi-objective mixed-integer linear programming (MOMILP) formulation: binary decision variables represent content-selection decisions, capacity constraints regulate the amount of recommended content, and hard policy constraints impose minimum or maximum requirements on welfare-related dimensions. At the same time, conflicting commercial and welfare-related objectives can be explored through weighted-sum or AUGMECON approaches, allowing the trade-offs between competing objectives to be made explicit and auditable.

The key challenge is therefore not whether recommendation systems can be optimized, but which objectives and constraints are incorporated into the allocation process. This distinction is particularly important in adolescent-oriented recommendation environments, where engagement-based objectives may conflict with behavioral-risk reduction, educational exposure, and content-diversity requirements. A constrained multi-objective allocation framework therefore provides a transparent way to examine these trade-offs without assuming a single welfare optimum.

2.5. Synthesis and research gap

Across these literatures, adolescent digital consumption may involve behavioral risks, engagement-based platform incentives may diverge from welfare-oriented objectives, and responsible recommendation provides mechanisms for incorporating broader objectives into allocation decisions. What remains underdeveloped is an auditable, ex-ante framework that jointly coordinates

DOI: 10.54254/2754-1169/2026.LD38497

commercial utility, behavioural risk, educational exposure, and content diversity under explicit operational constraints.

Motivated by this gap, the present study asks: How can a platform integrate welfare-oriented constraints into its recommendation decisions to balance commercial utility against behavioural risk, while ensuring educational value and content diversity? The paper answers this question by formulating recommendation as a constrained attention-allocation problem and developing a MOMILP that embeds modeled behavioural-risk control, educational exposure, and content diversity in the decision layer, shifting welfare considerations from ex-post correction to ex-ante allocation.

The purpose is not to claim real-world welfare effects from synthetic data, nor to identify a causal improvement in adolescent welfare. The risk term is a modeled behavioural-risk proxy, not a clinical or observational measure of harm. Through controlled synthetic experiments, the analysis quantifies feasible trade-offs among commercial utility, modeled risk proxy, educational exposure, and diversity, and examines their sensitivity and computational feasibility. It shows how governance constraints can be coordinated inside an allocation program. It does not show that such a program would causally raise welfare on live platforms.

3. Model

3.1. Economic environment and assumptions

This study models the video platform as an allocator of scarce recommendation exposure. For each adolescent user and candidate video, the platform decides whether the video is included in the recommendation set. Each assignment generates predicted short-term commercial value but may also contribute modeled behavioural risk and consume limited viewing-time, exposure, and delivery resources. The platform therefore maximizes governance-adjusted net value subject to user-level protection, category-diversity, item-exposure, and system-budget constraints. Selected user-level requirements are formulated as soft constraints with penalized slack variables to preserve feasibility when governance thresholds conflict.

3.2. Model ingredients: sets, parameters, and variables

Let I, J, and K denote the sets of adolescent users, candidate videos and content categories. For each $k \in K, C_k \subseteq J$ is the subset of videos in category k. The exogenous inputs v_{ij}, r_j, t_j, and c_j describe predicted commercial value, modeled behavioral risk, video duration, and allocation cost. To facilitate category-specific constrains, binary indicator variables δ_j^{ent} and δ_j^{edu} are defined to specify whether a video j belongs to Entertainment (Game/Relaxation) or Education category. Parameters T_i, R_i, K_i^{max}, K_i^{min}, A_i, E_i, M_k, N_j, and C_{total} are governance thresholds or capacities, while α, β, p_t, p_r, p_a, and p_e determine objective trade-offers and violation penalties. The binary variable x_{ij} equals one if video j is allocated to user i. Finally, ΔT_i^+, ΔR_i^+, ΔA_i^+, and ΔE_i^- measure the four soft-constraint violations. Complete definitions are provided in Appendix Table A1.

DOI: 10.54254/2754-1169/2026.LD38497

3.3. Experimental parameterization

The allocation inputs are synthetically generated to reproduce plausible differences across content categories rather than estimated from individual-level platform data. Governance thresholds, capacity limits, objective weights, and penalty coefficients define the baseline policy setting; the reported preference-weight sensitivity analysis varies α and β, *while the remaining policy parameters are held fixed.*

To reflect the development heterogeneity of adolescents, the user set is partitioned into two cohorts: junior high students (12-14 years old) and senior high students (15-18 years old). For the synthetic experiment, the two cohorts are assigned shares of 40% and 60%, respectively; this split is an education-stage-based modeling assumption informed by Educational Statistics in 2025. The video category is divided into four subcategories: Game, Relaxation, Lifestyle, and Education. For the synthetic experiment, this study assigns illustrative category shares of 30% Game, 40% Relaxation, 15% Lifestyle, and 15% Education. The video duration category is classified based on a video's duration t_j into three categories: Short ($t_j \leq 3min$), Medium (3min$< t_j \leq 15min$), Long ($t_j > 15min$).

Commercial value v_{ij} is generated from a category-, duration-, and age-specific truncated normal distribution; Table A.2c in Appendix A.2 reports the means, standard deviations, truncation bounds, and random seed. The modeled risk proxy r_j is formulated as product of video content risk and video length risk. Content-risk weights are assigned by category, and as a synthetic modeling assumption, shorter videos are assigned higher length-risk weights. The video duration t_j follows a log-normal distribution, simulating a long-tail ecosystem with a very large number of short videos and very few long videos. The value range is from 0.5 minutes to 60 minutes. The c_j is based on video duration using a linear model $c_j = \lambda \bullet t_j$, where the cost coefficient λ is assumed to be 0.001-unit cost per minute.

Viewing-time limits T_i are drawn from cohort-specific normal distributions, rounded to integer minutes, and clipped below at zero. R_i follows a stratified normal distribution across different age groups, and senior age group has greater risk tolerance than junior high students. This study set $K_i^{max} = floor(T_i/2)$, using a fixed reference duration of 2 minutes. And the minimum delivery requirement is set to $K_i^{min} = ceil(0.4K_i^{max})$. This is an illustrative policy calibration, and feasibility is verified for every user before optimization. A_i is defined as the floor of 60% of K_i^{max}. E_i is defined as the floor of 20% of K_i^{max}. The per-category exposure cap M_k is calibrated once to 9,275 and is held fixed in all MIP runs. It is not recomputed from the sample mean of $K_i^{\max}$ at solve time. This constant equals $0.35\times \mid I \mid \times 26.5 = 9{,}275$, using the instance size $\mid I \mid = 1{,}000$ and the rounded mean of $K_i^{\max}$ on this instance; the solver then treats $M_k = 9{,}275$ as a fixed right-hand side. N_j follows a Pareto distribution, where a small number of high-popularity videos are assigned extremely high thresholds, while most long-tail videos have thresholds of a few dozen. C_{total} is defined as 80% of the maximum possible total cost that is calculated by removing the cost constraint from the model.

The synthetic instance contains 1,000 users, 5,000 videos, and 5,000,000 potential user–video pairs, each associated with a simulated short-term commercial-value parameter. All synthetic data and solver programs can be found in https://github.com/elizaren/bcuaaw.

DOI: 10.54254/2754-1169/2026.LD38497

3.4. Objective, constraints, and properties

The optimization objective must balance two fundamentally competing goals, commercial value and cumulative modeled behavioural risk. Specifically, $\sum_{i\in I}\sum_{j\in J} v_{ij}x_{ij}$ captures predicted short-term commercial value, and $\sum_{i\in I}\sum_{j\in J} r_j x_{ij}$ internalizes the modeled risk indicator:

$$max\ Z_1 = \sum_{i\in I}\sum_{j\in J} v_{ij}x_{ij}$$

$$minZ_2 = \sum_{i\in I}\sum_{j\in J} r_j x_{ij}$$

Although the bi-objective formulation (Z_1, Z_2) conceptually clarifies the utility-risk trade-off, the real-world usage needs to balance theoretical rigor with deployability, this study therefore adopts the weighted-sum scalarization as our primary solver. This approach maintains deterministic convergence and low computational overhead. The scalarized objective integrates the maximization of commercial utility, the penalization of modeled risk proxy (via the negative coefficient $-\beta$ on Z_2), and feasibility penalties:

$$max\ Z_3 = \alpha \sum_{i\in I}\sum_{j\in J} v_{ij}x_{ij} - \beta\sum_{i\in I}\sum_{j\in J} r_j x_{ij} - \sum_{i\in I}\left(p_T\Delta T_i^+ + p_R\Delta R_i^+ + p_A\Delta A_i^+ + p_E\Delta E_i^-\right) \quad (1)$$

s.t.

$$\sum_{j\in J} x_{ij} \le K_i^{max},\ \ \forall i \in I \quad (2)$$

$$\sum_{j\in J} x_{ij} \ge K_i^{min},\ \ \forall i \in I \quad (3)$$

$$\sum_{j\in J} t_j x_{ij} \le T_i + \Delta T_i^+,\ \ \forall i \in I \quad (4)$$

$$\sum_{\mathrm{j\in J}} \mathrm{r_j x_{ij}} \le \mathrm{R_i} + \Delta \mathrm{R_i^+},\ \ \forall \mathrm{i} \in \mathrm{I} \quad (5)$$

$$\sum_{\mathrm{j\in J}} \delta_{\mathrm{j}}^{\mathrm{ent}} \mathrm{x_{ij}} \le \mathrm{A_i} + \Delta \mathrm{A_i^+},\ \ \forall \mathrm{i} \in \mathrm{I} \quad (6)$$

$$\sum_{\mathrm{j\in J}} \delta_{\mathrm{j}}^{\mathrm{edu}} \mathrm{x_{ij}} \ge \mathrm{E_i} - \Delta \mathrm{E_i^-},\ \ \forall \mathrm{i} \in \mathrm{I} \quad (7)$$

$$\sum_{\mathrm{i\in I}} \sum_{\mathrm{j\in C_k}} \mathrm{x_{ij}} \le \mathrm{M_k},\ \ \forall \mathrm{k} \in \mathrm{K} \quad (8)$$

$$\sum_{\mathrm{i\in I}} \mathrm{x_{ij}} \le \mathrm{N_j},\ \ \forall \mathrm{j} \in \mathrm{J} \quad (9)$$

$$\sum_{\mathrm{i\in I}} \sum_{\mathrm{j\in J}} \mathrm{c_j x_{ij}} \le \mathrm{C_{total}} \quad (10)$$

$$\mathrm{x_{ij}} \in \{0,1\},\ \ \forall \mathrm{i} \in \mathrm{I}, \forall \mathrm{j} \in \mathrm{J} \quad (11)$$

DOI: 10.54254/2754-1169/2026.LD38497

$$\Delta \mathrm{T}_{\mathrm{i}}^{+},\ \Delta \mathrm{R}_{\mathrm{i}}^{+} \geq 0,\ \ \forall \mathrm{i} \in \mathrm{I} \tag{12}$$

$$\Delta \mathrm{A}_{\mathrm{i}}^{+},\ \Delta \mathrm{E}_{\mathrm{i}}^{-} \in \mathbb{Z} \geq 0,\ \ \forall \mathrm{i} \in \mathrm{I} \tag{13}$$

p_T, p_R, p_A and p_E are penalty coefficients for the slack variables ΔT_i^+, ΔR_i^+, ΔA_i^+, and ΔE_i^-. Slack variables preserve feasibility when selected governance requirements conflict. The common penalty coefficient p = 0.01 discourages violations in the objective but does not guarantee their absence. Finally, α and β encode the modeled trade-off between short-term commercial utility and the aggregate modeled-risk indicator. Unless otherwise stated, the main comparative analysis uses α=0.834 and β=0.166 as the preference-weight setting; alternative weight combinations are examined in Section 5.3 sensitivity analysis.

The feasible set is defined by Constraints (2)-(13). Constraints (2)-(7) operate at the user level, Constrains (8)-(9) regulate concentration and Constraints (10)-(13) impose platform feasibility and variable domains.

Constraint (2) imposes a hard upper bound on the number of assigned videos to prevent 'choice overload' and decision fatigue. Constraint (3) imposes a per-user minimum delivery floor, preventing zero or very small recommendation lists. It does not, by itself, equalize aggregate delivery volume across methods. Constraint (4) limits total assigned duration. The slack ΔT_i^+ permits a measured violation when the duration threshold conflicts with other requirements; each unit is penalized in the objective. Constraint (5) suppresses high-risk content for adolescents. To further regulate content consumption, constraints are deployed on specific subcategories. Constraint (6) quantitatively curb the overexposure of purely entertainment-driven content. Conversely, constraint (7) enforces positive behavioural value guidance by retrieving at least E_i Education videos.

Constraint (8) enforces a strict upper limit M_k on the allocation of any single subcategory C_k. This constraint distributes allocated videos across multiple subcategories, thereby reducing excessive concentration within a single subcategory. Constraint (9) introduces a global maximum exposure quota N_j to prevent excessive concentration on a small subset of highly engaging videos. These constraints act as structural exposure regulators, they not only reduce popularity bias and winner-take-all, but also increase the opportunities of exploring high returns on long-tail videos. Additionally, constraint (10) imposes an operational cost limit to ensure that the allocation is practically economical.

Finally, the variable domains ensure logical consistency within the mixed-integer programming framework. The core assignment variables are binary, while the slack variables absorbing constraint violations are defined as non-negative continuous or integer values depending on their respective constraints.

The risk term appears twice by design, forming a hierarchical penalty structure. The global term $-\beta \sum r_j x_{ij}$ drives aggregate modelled risk reduction, while the per-user soft constraint $\sum_{j \in J} r_j x_{ij} \leq R_i + \Delta R_i^+$ with penalty p_R imposes a marginal cost on individual risk exceedance. With $P_R = 0.01 \ll \beta = 0.166$, the individual cap is deliberately weighted to be non-decisive; it does not enforce hard compliance but instead provides a soft discouragement against extreme user-level risk concentration. This design allows the optimizer to prioritize global gains while still penalizing—though not prohibiting—individual over-exposure.

DOI: 10.54254/2754-1169/2026.LD38497

4. Solution method

All experiments were implemented in Python and solved using Gurobi 13.0 on an Apple MacBook Pro M2 Max with 32 GB of RAM. Gurobi's standard mixed-integer optimization procedure, including presolve, linear-programming relaxations, cutting planes, and branch-and-bound search, was used to determine the binary user–video assignments.

Unless otherwise specified, the solver used a relative MIPGap of 1%, 12 threads, and no explicit time limit. The solver therefore terminated when the relative difference between the best feasible solution and the best available bound fell below 1%, meaning that the reported solutions were near-optimal within the specified tolerance rather than proven exact optima. The same dataset and solver settings were used for the proposed model and the baselines, while alternative preference weights and MIPGap settings were examined in the subsequent sensitivity and computational analyses.

This study uses an AUGMECON-inspired lexicographic ε-constraint procedure to explore the trade-off between penalized commercial value $P = Z_1 - p\sum \Delta$ and modeled risk Z_2, while reporting the original criteria (Z_1, Z_2). Lexicographic payoff gives $Z_2^{min} = 13{,}962.50$ and $Z_2^{nadir} = 96{,}736.10$ ($\max P = Z_1 - p\sum \Delta$, then min Z_2 on that face), so $r = 82{,}773.60$. With $q = 20$,

$$\varepsilon_k = 13{,}962.50 + \frac{k}{20} \times 82{,}773.60, \quad k = 0, \ldots, 20$$

For each grid point, the AUGMECON subproblem is

$$max\ Z_1 - p\sum \Delta + \rho\frac{s}{r}$$

subject to (2)-(13) and $Z_2 + s = \varepsilon_k,\ s \geq 0$, with $\rho = 10^{-3}$.

To prioritize penalized commercial value before surplus maximization, this study uses a two-stage lexicographic implementation. Stage A maximizes P; Stage B maximizes s subject to $P \geq P_A^{inc} - 10^{-4}$, where P_A^{inc} is the Stage-A incumbent. Both stages use a relative MIPGap of 0.1%.

5. Result

The experiments evaluate the proposed allocation model along four dimensions: comparison with a myopic commercial baseline, sensitivity to the objective weights and an AUGMECON approximation of the value-risk frontier, computational performance under alternative MIPGap settings, the diversity of the resulting exposure allocation. All comparisons use the same synthetic input data and common operational constraints unless otherwise stated. The analysis therefore evaluates model behaviour under controlled assumptions; it does not estimate causal effects on adolescent welfare.

DOI: 10.54254/2754-1169/2026.LD38497

5.1. Baselines

The evaluation compares the proposed framework against two baselines. The purpose is not to reproduce the full complexity of industrial recommendation systems, but to isolate allocation logics that dominate current practice and build a foundation for evaluating the proposed model: (1) engagement-driven optimization without welfare constraints, and (2) the same engagement-driven slate followed by ex-post hard safeguards. All methods use the same synthetic dataset. The proposed model and Baseline 1 are solved with a 1% MIPGap, whereas Baseline 2 is a deterministic post-processing procedure.

5.1.1. Baseline 1: myopic commercial value maximization

The first baseline is designed to maximizes predicted commercial value alone, using the same exogenous inputs and capacity parameters as the proposed model. It omits the risk term, slack penalties, entertainment/education bounds, and category caps, while retaining the common minimum and maximum delivery bounds, the global maximum exposure quota, and the operational cost limit. Especially, the baseline 1 uses a hard viewing-time constraint without the slack variable.

$$\max Z_4 = \sum_{i\in I}\sum_{j\in J} v_{ij}x_{ij} \tag{14}$$

s.t.

$$\sum_{j\in J} x_{ij} \le K_i^{max}, \quad \forall i \in I \tag{2}$$

$$\sum_{j\in J} x_{ij} \ge K_i^{min}, \quad \forall i \in I \tag{3}$$

$$\sum_{j\in J} t_j x_{ij} \le T_i, \quad \forall i \in I \tag{15}$$

$$\sum_{i\in I} x_{ij} \le N_j, \quad \forall j \in J \tag{9}$$

$$\sum_{i\in I}\sum_{j\in J} c_j x_{ij} \le C_{total} \tag{10}$$

$$x_{ij} \in \{0,1\}, \quad \forall i \in I, \forall j \in J \tag{11}$$

5.1.2. Baseline 2: myopic allocation with post-hoc hard filtering

The second baseline captures the ex-post safeguard logic that motivates this paper. This study first keeps the baseline 1 solution unchanged—so the allocation is still chosen only to maximize predicted commercial value, this study then attempts to repair the slate user by user by enforcing the hard risk and entertainment bounds and adding educational items where feasible; any residual educational shortfalls are reported explicitly:

$$\sum_{j\in J} r_j x_{ij} \le R_i, \quad \forall i \in I \tag{16}$$

$$\sum_{j\in J} \delta_j^{ent} x_{ij} \le A_i, \quad \forall i \in I \tag{17}$$

$$\sum_{j\in J} \delta_j^{edu} x_{ij} \ge E_i, \quad \forall i \in I \tag{18}$$

DOI: 10.54254/2754-1169/2026.LD38497

Violations of the modelled risk or Entertainment upper bounds are removed in descending order of $r_j/(v_{ij}+\mathscr{E})$; Educational shortfalls are filled, when feasible, by adding Education videos in descending order of v_{ij}. Residual Education shortfalls are reported rather than relaxed—consistent with hard ex-post interception in practice. Finally, feasible residual arcs are refilled in pure value order, first targeting the proposed model's realized cost and then its delivery volume. This procedure deliberately does not re-optimize $\alpha\sum_{i\in I}\sum_{j\in J}v_{ij}x_{ij}-\beta\sum_{i\in I}\sum_{j\in J}r_jx_{ij}$; both procedures start from the same value scores, the comparison isolates the consequences of joint optimization versus the specified post-hoc repair rule, subject to the important differences in achieved volume, feasibility, and slack treatment reported below.

5.2. Comparative analysis

The evaluation was not restricted to a single objective score. Instead, four groups of evidence were assembled: comparative performance against two baselines, sensitivity to α and β, computational performance under different MIPGap settings, and content-ecology diversity analysis. This design aims to answer not only whether the proposed model works, but also how stable, expensive, and structurally balanced the resulting solution.

Table 1 presents the three-way comparison. Under a nearly identical delivery cost (42.38 vs. 42.75), the proposed model reduces modelled risk by 58.9% relative to the baseline 1 while retaining 68.5% of its predicted commercial value. Delivery volume falls by 21.1%, yet aggregate assignment time rises slightly (+0.87%), indicating that risk reduction is not achieved by merely shrinking exposure. Decomposing modeled risk as $R=N\bar{r}$, where N is the number of deliveries and $\bar{r}$ the mean risk per delivery. As Table 2 shown, the cut from $N_b=26{,}428$ to $N_p=20{,}851$ accounts for 35.8% of the total risk decrease when myopic baseline 1 $\bar{r}$ is fixed. Under a baseline-profile counterfactual decomposition, the delivery-count reduction accounts for 35.8% of the total decrease, while the remaining 64.2% is associated with a lower mean modeled risk per delivery ($\bar{r}$ decreases from 4.83 to 2.51). Equivalently, a volume-only counterfactual that preserves the myopic baseline risk profile would leave modelled risk near 100,619—almost twice the proposed level. Thus, the governance constraints operate mainly by reallocating the slate under an almost unchanged time/cost budget.

Table 1. Three-way comparison at α=0.834, β=0.166, P=0.01

Metric	Baseline 1	Baseline 2	Proposed
predicted commercial value	26,341.23	8,182.46	18,044.28
Modelled risk	127,531.35	28,896.95	52,356.00
Delivery cost	42.38	29.81	42.75
Delivery volume	26,428	8,478	20,851
Value retained vs. myopic	100%	31.1%	68.5%
Risk change vs. myopic	-	−77.3%	−58.9%
Cost gap vs. proposed target	−0.37	−12.93	0
Volume gap vs. proposed target	+5,577	−12,373	0

DOI: 10.54254/2754-1169/2026.LD38497

Table 1. (continued)

Users violating hard R_i / A_i	-	0 / 0	via ΔR_i^+ , ΔA_i^+
Users missing hard E_i	-	902 / 1,000	via ΔE_i^-
Videos dropped / education topped up	-	19,722 / 1,722	-

Table 2. Decomposition of the modelled-risk reduction: volume versus structure

Source of ΔRisk	Contribution	Share of risk reduction
Delivery-volume reduction	−26,912	35.8%
Delivery-structure change	−48,263	64.2%

Note. Total modeled risk is decomposed as $R = N\bar{r}$, where N is the number of deliveries and $\bar{r}$ the mean risk per delivery. Let subscripts b and p denote the baseline 1 and proposed solutions, respectively. The change in modeled risk is $\Delta R = R_p - R_b = N_p\bar{r}_p - N_b\bar{r}_b$. Under a baseline-profile counterfactual decomposition, the volume component is $\Delta R_{\text{vol}} = (N_p - N_b)\bar{r}_b$ and the structure component is $\Delta R_{\text{struct}} = N_p(\bar{r}_p - \bar{r}_b)$, with $\Delta R = \Delta R_{\text{vol}} + \Delta R_{\text{struct}}$. Shares are reported as $|\Delta R_{\text{vol}}| / |\Delta R|$ and $|\Delta R_{\text{struct}}| / |\Delta R|$.

The baseline 2 is the critical test of the paper's motivating claim. Enforcing the same R_i, A_i, and E_i after myopic allocation does reduce 77.3% of the modeled risk and eliminates hard risk/Entertainment violations. However, it does so by discarding 19,722 assignments and decreases predicted commercial value to 31.1% of the baseline 1 level. Education top-ups recover only 1,722 items; 902 of 1,000 users still miss the Educational floor under hard filtering. Value-ordered refill then adds only 50 further assignments: once high-value (typically high-risk) candidates are admitted, residual risk budgets are exhausted for almost all users, so this value-ordered refill procedure did not reach the proposed model's cost or delivery volume under the imposed hard limits (cost gap −12.93, volume gap −12,373). The specified repair procedure did not recover the proposed model's cost or delivery volume. This is a limitation diagnostic for that repair rule, not proof that every hard-constraint allocation is infeasible at matched volume or cost.

Two implications follow. First, ex-post hard safeguards change the feasible set after the fact but do not rewrite initial allocation logic; the resulting slate is a truncated engagement ranking, not a welfare-aware optimum. Second, the proposed model is not merely 'myopic plus constraints.' By internalising modeled risk in the objective and absorbing residual conflicts through penalised slacks —while keeping delivery cost aligned with the myopic budget—it retains substantially more predicted commercial value than post-hoc filtering (18,044 vs. 8,182) at a moderate risk level, the proposed model permits penalized educational shortfalls, so its education compliance must be assessed from the realized slack values. In short, embedding modeled-risk and governance terms in the decision layer yields a different allocation from applying the specified hard repair rule ex post; however, the two realized solutions differ in volume and constraint satisfaction.

From a governance perspective, the proposed model changes the allocation criterion from commercial-value maximization to optimization under explicit modeled-risk, educational, and diversity considerations. The resulting reduction in predicted commercial value therefore represents the modeled opportunity cost of incorporating these governance objectives. Quantifying whether this trade-off improves long-term welfare requires behavioral and longitudinal data beyond the current synthetic experiment.

DOI: 10.54254/2754-1169/2026.LD38497

5.3. Sensitivity analysis

As shown in Table 3, along the tested weight path from $(\alpha\beta) = (1.00.0)$ to $(0.01.0)$, predicted commercial value decreases from $25{,}109.70$ to $3{,}776.98$, while the modeled-risk score decreases from $99{,}773.75$ to $13{,}968.50$. The response is largely monotone, but a narrow preference cliff appears at $\alpha \in [0.8320.8345]$. This study selects $\alpha = 0.834$, $\beta = 0.166$ as an illustrative operating point at the onset of the cliff. On the AUGMECON grid in Section 5.4, it lies between $k = 9$ ($Z_1 = 17{,}844.56$, $Z_2 = 51{,}204$) and $k = 10$ ($Z_1 = 18{,}702.62$, $Z_2 = 55{,}268$); the nearest grid point in modeled risk is $k = 9$. This pattern is illustrated in Figure 1. This study therefore uses this weight combination as a representative setting for comparison with baselines. This selection is specific to the evaluated weight grid and does not imply global optimality. Overall, the weighted-sum path is consistent with diminishing marginal value gains as additional emphasis is placed on commercial utility. The sensitivity analysis characterizes the behavior of the weighted-sum solver; however, this study next benchmark the selected point against the valuated trade-off grid in section 5.4.

Table 3. Model outputs under different weight settings (p=0.01)

α	β	Commercial Value	Modeled Risk
1.0000	0.0000	25,109.70	99,773.75
0.9000	0.1000	24,087.65	83,688.00
0.8500	0.1500	22,384.31	73,498.00
0.8400	0.1600	20,671.64	64,752.00
0.8350	0.1650	19,413.91	58,781.00
0.8345	0.1655	19,184.70	57,573.00
0.8340	0.1660	18,044.28	52,356.00
0.8335	0.1665	17,307.79	48,885.00
0.8330	0.1670	16,587.93	45,510.00
0.8320	0.1680	15,392.11	39,793.00
0.8310	0.1690	14,921.36	37,609.00
0.8300	0.1700	14,760.44	36,888.50
0.8200	0.1800	13,808.10	32,660.00
0.8000	0.2000	12,441.80	27,043.50
0.6000	0.4000	8,626.36	14,302.10
0.4000	0.6000	8,443.27	14,118.30
0.2000	0.8000	8,141.35	13,995.00
0.0000	1.0000	3,776.98	13,968.50

DOI: 10.54254/2754-1169/2026.LD38497

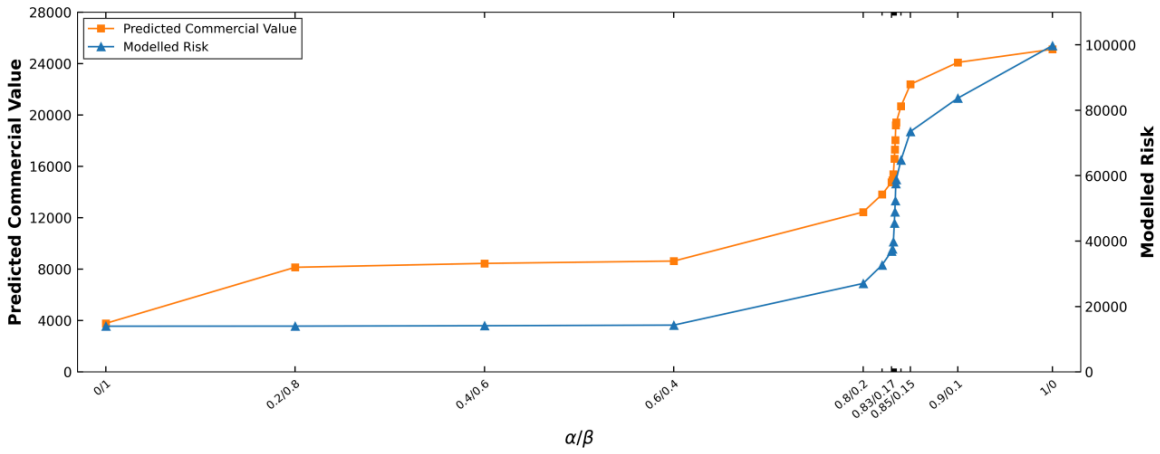


Figure 1. Sensitivity analysis with respect to different α/β weights

The curved trade-off illustrates that alternative preference weights generate materially different commercial-value–risk combinations. Selecting an operating point therefore requires an explicit stakeholder-preference rule; future work could elicit such preferences and refine the evaluated grid.

5.4. AUGMECON approximation of the value–risk trade-off

An AUGMECON grid with $q = 20$ approximates the value–risk trade-off, following Section 4. The lexicographic payoff range $[13{,}962.50\ 96{,}736.10]$ is partitioned into 20 steps of 4,138.68. For each ε_k this study solves the two-stage subproblem. All 21 subproblems returned Gurobi OPTIMAL status within the specified tolerances; the reported Stage-A gaps ranged from 0.007% to 0.092%. Table 4 and Figure 2 report integer-feasible criteria. Surplus s_k is the Stage-B surplus in $Z_2 + s = \varepsilon_k$, not leftover budget from an inequality. No bypass occurred.These points are used to characterize the trade-off within the evaluated grid; linearly interpolated points between integer-feasible solutions are not treated as feasible solutions or as a basis for calculating an optimization gap.

Relative to the selected weighted-sum solution ($Z_1 = 18{,}044.28$, $Z_2 = 52{,}356$), the nearest AUGMECON vertex is $k = 9$ ($Z_1 = 17{,}844.56$, $Z_2 = 51{,}204$). The weighted point therefore has slightly higher commercial value at higher modelled risk; it is not claimed to coincide with $k = 9$, and no interpolation along the grid is used as a certified gap.

Table 4. AUGMECON grid of the value–risk trade-off (q=20, p=0.01). Z1* and Z2* are the original criteria after the lexicographic subproblem; sk is the optimized surplus in Z2+s=εk. Time is Stage A + Stage B. Rows follow the computational order from the loosest to the tightest risk cap

k	ϵ_k	Z_2^*	Z_1^*	s_k	Time (s)	Stage-A gap (%)	Stage-B status
20	96,736.10	95,682.55	25,060.03	1,053.55	585.46	0.086	Optimal
19	92,597.42	92,296.75	24,891.30	300.67	458.79	0.075	Optimal
18	88,458.74	88,291.00	24,561.36	167.74	520.83	0.053	Optimal
17	84,320.06	84,169.25	24,172.20	150.81	525.44	0.069	Optimal
16	80,181.38	80,096.75	23,589.24	84.63	434.11	0.056	Optimal

DOI: 10.54254/2754-1169/2026.LD38497

Table 4. (continued)

15	76,042.70	75,995.25	22,862.46	47.45	418.99	0.038	Optimal
14	71,904.02	71,833.50	22,076.04	70.52	476.81	0.061	Optimal
13	67,765.34	67,690.50	21,265.64	74.84	357.30	0.069	Optimal
12	63,626.66	63,535.50	20,428.50	91.16	498.58	0.091	Optimal
11	59,487.98	59,413.50	19,574.24	74.48	374.94	0.080	Optimal
10	55,349.30	55,268.00	18,702.62	81.30	579.61	0.092	Optimal
9	51,210.62	51,204.00	17,844.56	6.62	527.42	0.012	Optimal
8	47,071.94	47,057.00	16,965.45	14.94	388.93	0.022	Optimal
7	42,933.26	42,924.00	16,085.77	9.26	414.55	0.017	Optimal
6	38,794.58	38,780.50	15,199.68	14.08	397.52	0.025	Optimal
5	34,655.90	34,645.00	14,297.84	10.90	505.82	0.024	Optimal
4	30,517.22	30,505.00	13,336.01	12.22	493.67	0.030	Optimal
3	26,378.54	26,373.00	12,300.06	5.54	301.34	0.020	Optimal
2	22,239.86	22,208.30	11,163.74	31.56	348.03	0.087	Optimal
1	18,101.18	18,097.10	9,942.73	4.08	406.91	0.022	Optimal
0	13,962.50	13,962.50	7,935.54	0.00	215.43	0.007	Optimal

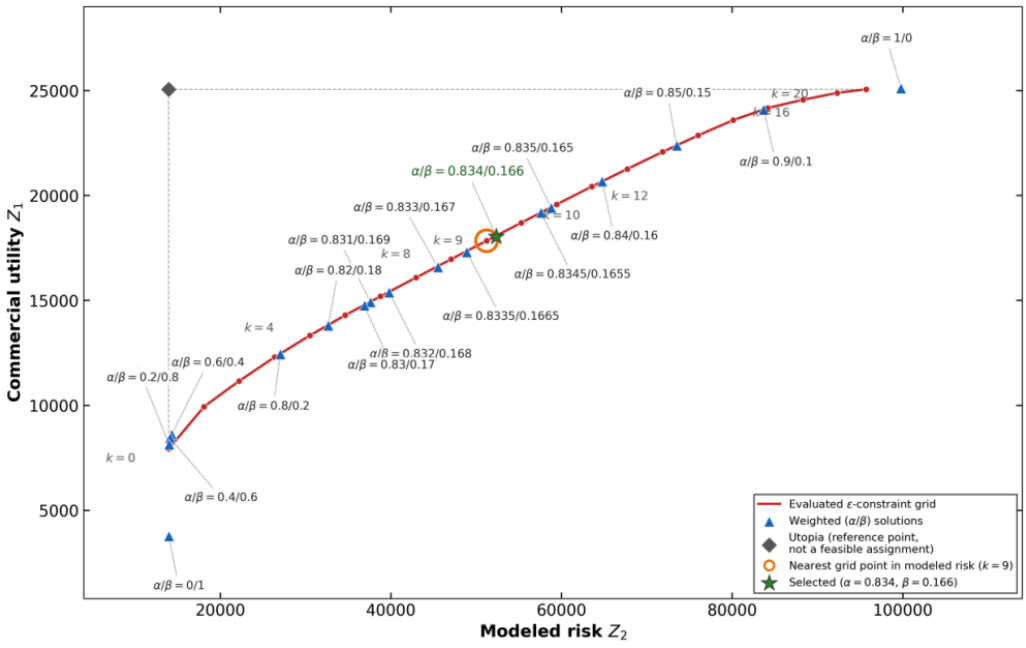


Figure 2. AUGMECON grid approximation of the value–risk trade-off and the weighted-sum solutions

DOI: 10.54254/2754-1169/2026.LD38497

5.5. Computational performance

As Table 5 shown, in the reported run on the stated hardware, tightening the MIPGap from 1% to 0.01% increased runtime from 67.80 s to 185.53 s, while peak RAM use remained approximately unchanged and the objective increased by 0.59%. Because reducing the MIPGap tolerance below 1% yielded only small improvements in the best feasible objective value while requiring longer solution times, a target MIPGap of 1% was used in the main experiments.

Table 5. Solver outputs under different MIPGap settings (α=0.834, β=0.166)

MIPGap	Incumbent (Objective)	Best Bound	Actual Gap (%)	Runtime (s)	RAM (GB)	Terminated?
1%	6,116.48	6,153.38	0.603354	67.80	22.68	Optimal
0.1%	6,147.41	6,153.38	0.0971238	68.69	22.68	Optimal
0.01%	6,152.46	6,153.07	0.0099766	185.53	22.54	Optimal

5.6. Diversity analysis and discussion

Figure 3 compares recommendation volume by content category under the baseline 1 and proposed models. Table 6 reports category-composition and user-level diversity metrics; entropy is the pooled Shannon entropy of aggregate category shares (natural logarithm, in nats), not the average of per-user entropies. Relative to the baseline 1, the proposed model raises pooled category entropy and within-user genre coverage, while video-level concentration increases (fewer unique videos, higher Gini and Top-100 share).

Homogenization is reduced but not eliminated: education still accounts for 44.24% of pushes, and the P90 max-category share remains 70.04%. The educational floor E_i thus replaces engagement-driven homogenization with softer policy-driven concentration. The global cap $\sum_{i\in I}\sum_{j\in C_k} x_{ij} \le M_k,\ \forall k \in K\ (M_k = 9275)$ is nearly binding for Education (9,225; slack 50) but does not limit within-user shares, so list-level dominance can persist.

Two limitations follow. The model has no participation/abandonment term, so the analysis does not capture how engagement may change under a more educational content mix. The reported value-retention rate is based on synthetic commercial-value scores and is not a forecast of realized platform revenue.

Table 6. reports category-composition and user-level diversity metrics

Indicator	Baseline 1	Proposed	Δ
Aggregate category composition			
Pooled Category entropy (nats)	0.612	1.174	+0.562
Unique videos	3,375	1,937	-1,438
Video push Gini	0.364	0.464	+0.100

DOI: 10.54254/2754-1169/2026.LD38497

Table 6. (continued)

Top-100 video share	7.50%	18.62%	+11.12 pp
User-level diversity			
Avg. distinct categories / user	2.04	3.62	+1.58
Avg. max-category share / user	71.13%	55.14%	-16.00 pp

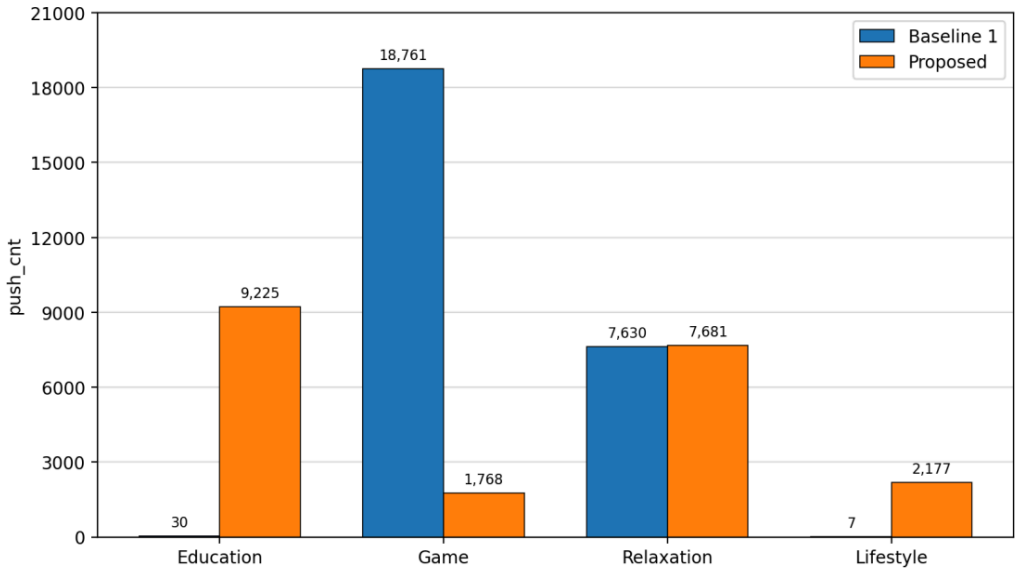


Figure 3. Recommendation counts by category (baseline 1 vs proposed)

6. Conclusion

This study develops a mixed-integer allocation model that balances predicted commercial value against a synthetic behavioural-risk proxy while imposing education, diversity, exposure, and budget constraints. On the reported synthetic instance, the selected specification reduced modeled risk by 58.9% relative to Baseline 1 while retaining 68.5% of predicted commercial value at a nearly identical delivery cost. Compared with post-hoc repair, joint optimization produces a different allocation, although the comparison is not an equal-volume test. These findings illustrate how governance objectives can be incorporated directly into recommendation allocation, but they do not establish causal welfare gains or real-world platform performance. The results depend on synthetic parameters, a synthetic commercial-value proxy, and a specific preference-weight setting; large-scale deployment may also require approximate or decomposition-based methods. Future work should validate the parameterization with platform and behavioural data, elicit stakeholder preferences, and evaluate matched-volume comparators and scalable solution methods.

References

[1] Xu, Z., Gao, X., Wei, J., Liu, H., & Zhang, Y. (2023). Adolescent user behaviors on short video application, cognitive functioning and academic performance. *Computers and Education, 203*, Art. no. 104865. doi:

DOI: 10.54254/2754-1169/2026.LD38497

10.1016/j.compedu.2023.104865.
[2] Keles, B., McCrae, N., & Grealish, A. (2020). A systematic review: The influence of social media on depression, anxiety and psychological distress in adolescents. *International Journal of Adolescence and Youth, 25*(1), 79–93. doi: 10.1080/02673843.2019.1590851.
[3] Allcott, H., Gentzkow, M., & Song, L. (2022). Digital addiction. *American Economic Review, 112*(7), 2424–2463. doi: 10.1257/aer.20210867.
[4] Orben, A., & Przybylski, A. K. (2019). The association between adolescent well-being and digital technology use. *Nature Human Behaviour, 3*(2), 173–182. doi: 10.1038/s41562-018-0506-1.
[5] Odgers, C. L., & Jensen, M. R. (2020). Annual research review: Adolescent mental health in the digital age: Facts, fears, and future directions. *Journal of Child Psychology and Psychiatry, 61*(3), 336–348. doi: 10.1111/jcpp.13190.
[6] Valkenburg, P. M., Meier, A., & Beyens, I. (2022). Social media use and its impact on adolescent mental health: An umbrella review of the evidence. *Current Opinion in Psychology, 44*, 58–68. doi: 10.1016/j.copsyc.2021.08.017.
[7] Hébert, B., & Zhong, W. (2026). Engagement maximization. *Theoretical Economics, 21*(3), 849–887. https://doi.org/10.3982/te6567
[8] Allcott, H., Braghieri, L., Eichmeyer, S., & Gentzkow, M. (2020). The welfare effects of social media. *American Economic Review, 110*(3), 629–676. https://doi.org/10.1257/aer.20190658
[9] Calvano, E., Calzolari, G., Denicolò, V., & Pastorello, S. (2024). Artificial intelligence recommendations: Evidence, issues, and policy. *Oxford Review of Economic Policy, 40*(4), 843–853. https://doi.org/10.1093/oxrep/grae048
[10] Heidhues, P., Köster, M., & Kőszegi, B. (2023). Steering fallible consumers. *Economic Journal, 133*(652), 1430–1465. https://doi.org/10.1093/ej/ueac093
[11] Milano, S., Taddeo, M., & Floridi, L. (2020). Recommender systems and their ethical challenges. *AI & Society, 35*(4), 957–967. https://doi.org/10.1007/s00146-020-00950-y
[12] Stray, J., et al. (2024). Building human values into recommender systems: An interdisciplinary synthesis. *ACM Transactions on Recommender Systems, 2*(3), Art. no. 20, 1–57. https://doi.org/10.1145/3632297
[13] Helberger, N., Karppinen, K., & D'Acunto, L. (2018). Exposure diversity as a design principle for recommender systems. *Information, Communication & Society, 21*(2), 191–207. doi: 10.1080/1369118X.2016.1271900.
[14] Mladenov, M., et al. (2020). Optimizing long-term social welfare in recommender systems: A constrained matching approach. In *Proceedings of the 37th International Conference on Machine Learning (ICML) (Volume 119)* (pp. 6987–6998).
[15] Jannach, D., & Abdollahpouri, H. (2023). A survey on multi-objective recommender systems. *Frontiers in Big Data, 6*, Art. no. 1157899. doi: 10.3389/fdata.2023.1157899.
[16] Zheng, Y., & Wang, D. X. (2022). A survey of recommender systems with multi-objective optimization. *Neurocomputing, 474*, 141–153. doi: 10.1016/j.neucom.2021.11.041.
[17] Ehrgott, M. (2005). *Multicriteria optimization* (2nd ed.). Springer. doi: 10.1007/3-540-27659-9.
[18] Mavrotas, G. (2009). Effective implementation of the ε-constraint method in multi-objective mathematical programming problems. *Applied Mathematics and Computation, 213*(2), 455–465. doi: 10.1016/j.amc.2009.03.037.