
Unveiling User Satisfaction and Creator Productivity Trade-Offs in Recommendation Platforms

Fan Yao*
University of Virginia
fy4bc@virginia.edu

Yiming Liao
Meta
yimingliao@meta.com

Jingzhou Liu
Meta
jingzhol@meta.com

Shaoliang Nie
Meta
snie@meta.com

Qifan Wang
Meta
wqfcr@meta.com

Haifeng Xu
University of Chicago
haifengxu@uchicago.edu

Hongning Wang
University of Virginia
hw5x@virginia.edu

Abstract

On User-Generated Content (UGC) platforms, recommendation algorithms significantly impact creators’ motivation to produce content as they compete for algorithmically allocated user traffic. This phenomenon subtly shapes the volume and diversity of the content pool, which is crucial for the platform’s sustainability. In this work, we demonstrate, both theoretically and empirically, that a purely relevance-driven policy with low exploration strength boosts short-term user satisfaction but undermines the long-term richness of the content pool. In contrast, a more aggressive exploration policy may slightly compromise user satisfaction but promote higher content creation volume. Our findings reveal a fundamental trade-off between immediate user satisfaction and overall content production on UGC platforms. Building on this finding, we propose an efficient optimization method to identify the optimal exploration strength, balancing user and creator engagement. Our model can serve as a pre-deployment audit tool for recommendation algorithms on UGC platforms, helping to align their immediate objectives with sustainable, long-term goals.

1 Introduction

User-generated content (UGC) platforms have become an indispensable component of our daily lives [5, 45]. Those platforms, including various social media (e.g., Facebook, Instagram), streaming services (e.g., YouTube, TikTok) and many more, count on algorithmic recommendation algorithms [36, 5] to help content consumers (i.e., users) navigate the vast ocean of content generated by creators. Unlike other content recommendation platforms such as Netflix and Spotify, user experience on UGC platforms critically relies on the active participation of creators [58], as the goal of enhancing user engagement is inherently linked to the abundance and diversity of the content pool.

Recent studies have begun to explore how a platform’s algorithmic decisions, such as their employed recommendation algorithms and revenue sharing agreements, might influence the behavior of content creators and subsequently affect user welfare [52, 54, 55, 34, 30, 33, 32, 43]. A common technique used in these works is to model the competition among creators who strive to establish their brand and comparative advantages by selecting topics that maximize the traffic or rewards from the platform. While these competition models offer valuable insights into how platforms’ interventions could lead to suboptimal outcomes in terms of content diversity and issues related to popularity bias, they ignore the dimension of content creation volume — an equally critical aspect of such competition

*Research conducted during an internship at Meta. Corresponding author.

dynamics. In fact, there are evidence suggesting that traffic received directly influences creator productivity [31, 56]. A recent study on a leading video recommendation platform found that creators whose content received boosted exposure significantly increased their video production without compromising quality [31]. And a recent field experiment on a large-scale video sharing social network showed that while popularity-based recommendation strategies boost content consumption, they can reduce content production [56]. The similar effect has also been documented on Instagram. For instance, when the recommender system allocates more traffic to targeted groups of creators, their production frequency increases, as indicated by metrics such as the number of daily active creators and the daily average creation volume per group. However, disproportionately boosting traffic to certain creator groups can negatively impact others; for instance, directing more traffic to popular, or “head”, creators often diminishes engagement from less prominent, or “tail”, creators. And overall, a statistically significant positive correlation exists between content viewership and the corresponding creator’s productivity².

Motivated by real-world evidence and the research gap in modeling production frequency within content creation competition, we introduce a new game-theoretical model, named Cournot Content Creation Competition (C^4), aiming to study the impact of a platform’s recommendation strategy on creators’ production willingness. Our C^4 framework builds upon the Content Creator Competition (C^3) framework introduced in [52, 54, 55], assuming that creators compete for platform-allocated user traffic [27, 29]. However, unlike C^3 and similar previous models, our approach models the competition where creators are aware of their expertise and consistently produce within their niche, but strategically adjust their production frequency to balance gain and cost, which is often observed on mature platforms such as YouTube and Instagram. Our proposed competition model resonates with the well-established Cournot competition [14] in economics, where firms compete for revenue by strategically setting their production quantities. Hence, it inherits its name.

Our C^4 framework offers a powerful tool for analyzing the competition dynamics among creators, as it always yields a unique Pure Nash equilibrium (PNE) that enables a precise prediction of the total content creation volume under any specific recommendation strategy. Furthermore, our in-depth analysis of C^4 ’s equilibrium reveals a critical and interesting insight: while increased recommendation accuracy boosts immediate user satisfaction, it simultaneously reduces creators’ motivation to produce content, potentially compromising long-term user engagement. This finding, supported by both theoretical analyses and simulations, suggests the necessity to balance user and creator engagement through a careful control over the recommendation algorithm’s exploration strength at a per-user basis. We formulated this mechanism design challenge as a bi-level optimization problem and tackled it using a projected gradient descent approach with an efficient gradient approximation scheme, providing an effective method to achieve the optimal trade-off between user satisfaction and creator productivity.

In summary, our contributions are threefold: (a) **modeling-wise**, we introduce a new game-theoretical framework, C^4 , to investigate how recommendation algorithms affect content creation frequency among creators; (b) **conceptually**, we reveal a new insight that managing the exploration strength of the recommendation algorithm can balance between short-term user satisfaction and long-term creator engagement at equilibrium; and (c) **technique-wise**, we reformulate the mechanism design problem of identifying the optimal engagement trade-off at the equilibrium into a solvable offline optimization problem, tackled using approximated gradient descent. Our C^4 framework and its derived solution serve as a pre-deployment audit tool for platforms, assessing the effects of algorithmic choices on creator and user engagement.

2 Related Work

The study of online content creation economy has captured the attention of machine learning community recently, leading to a diverse collection of models addressing the dynamics of content creator competition [25, 24, 3, 54, 52, 57, 32, 34, 30, 16]. In these models, creators strategically select the type [49], topic [30, 34, 52], or quality [25, 32] of their content, competing for resources such as traffic [30, 2, 24], user engagement [52], or platform-provided incentives [57, 54]. Some models aim to explore the properties of creator-side equilibrium, investigating how creators specialize at

²Further details regarding the Instagram findings are protected under a non-disclosure agreement (NDA) and may be disclosed in future versions of this work.

equilibrium [34], the impact of creators’ strategic behaviors on social welfare [52], and the design of optimization methods for long-term welfare considering these behaviors [2, 3, 54, 57, 32, 33, 39]. Unlike these works, which generally assume equally paced creation frequency among creators, our proposed C^4 games consider scenarios where creators strategically control their creation quantities, allowing us to analyze how recommendation algorithms influence overall content creation volume at the equilibrium.

Our model of the platform’s recommendation algorithm draws inspiration from the proportional allocation concept in game theory, applicable to resource distribution [9, 41] and contest design [48]. We assume that each user contributes a unit of traffic, which is allocated to creators based on both of their merit and effort (content creation frequency). This modeling is closely related to the Tullock contest [48], also known as the lottery contest, where the probability of winning a fixed prize is proportionate to the effort expended relative to the total effort by all contestants. While the Nash equilibrium of the one-dimensional Tullock contest with homogeneous costs is well-understood [20, 21], our work extends this framework to include heterogeneous convex costs. A recent study [53] also explored competition between human and Generative AI creators within a similar setup, examining its impact on total content creation volume. However, they did not address the influence of different recommendation algorithms, which we investigate in our work.

In a broader sense, our work contributes to a line of research evaluating the impact of recommender systems on individuals, specifically exploring how deployed algorithms shape user and content creator behavior [16, 12, 11, 38, 17, 35, 19] and how we can design new algorithms to address these effects [10, 50, 6, 51, 4, 43, 1]. Our study introduces a key insight: recommender algorithms optimized solely for user satisfaction can unintentionally reduce content creators’ willingness to engage, thereby impacting long-term user engagement. We address this challenge by proposing a solution that balances creator engagement and user satisfaction through imposing exploration strengths tailored to individual users.

3 The Cournot Content Creation Competition

In this section, we introduce the formulation of Cournot Content Creation Competition (C^4), which models the competition among creators for traffic on UGC platforms. This model considers the potential impact of the platform’s traffic reallocation mechanisms, such as their deployed recommendation algorithms, where creators strategically choose their production frequencies to optimize their allocated traffic. Each C^4 game instance \mathcal{G} is characterized by a tuple $(n, m, M, \{c_i\}_{i=1}^n, \{\beta_i\}_{i=1}^m)$. We detail each component of this tuple as follows:

1. **Basic setups.** There is a set of content creators denoted by $[n] = \{1, \dots, n\}$, and a set of users denoted by $\{u_j\}_{j=1}^m$. We assume each user j has a stable preference over creators and such relationship is captured by an n -by- m matrix M with its (i, j) -th entry $w_{ij} \in [0, 1]$ denoting the strength of user j ’s preference over creator i ’s content. Each creator determines a production frequency $x_i \in \mathbb{R}_{\geq 0}$ in a unit amount of time (e.g., one week/one month). For the purpose of our analysis, x_i can be interpreted interchangeably as either the production frequency or volume, provided there is no ambiguity³. Follow the terminology of game theory literature, x_i is referred to as the action or pure strategy of creator i . Each creator i is associated with a cost function c_i , which characterizes the cost for creating content at frequency x_i . We have two assumptions about c_i : 1. c_i is increasing in x_i and $c_i(x_i) \rightarrow +\infty$ as $x_i \rightarrow +\infty$, which reflects that content creation is always not free. 2. c_i is convex in x_i , indicating a non-decreasing marginal cost of improving production frequency.
2. **Platform intervention.** We assume that each user u_j contributes a unit amount of traffic, and the platform redistributes the total user traffic based on a relevance-based recommendation algorithm that adheres to certain probabilistic principle. Specifically, the recommendation algorithm matches u_j to each piece of content produced by creator i with a probability proportional to $\exp(\beta_j \sigma_{ij})$, where $\sigma_{ij} = w_{ij} + \epsilon_{ij}$ represents the algorithm’s estimated relevance score. This score combines the true preference score w_{ij} with independent Gaussian noise $\epsilon_{ij} \sim \mathcal{N}(0, \sigma^2)$. The parameter $\beta_j \geq 0$ governs the exploration of the matching for user u_j : a higher β_j results in a more precise

³For the elegance of our theoretical analysis, we treat x_i as a continuous variable, although the key messages and insights of this paper are preserved if x_i is discrete.

matching, while a lower β_j introduces more exploration into the matching results ⁴. In this work, we analyze intervention mechanisms that operate under this Personalized Probabilistic Matching (PPM) principle, parameterized by $\beta = \{\beta_i\}_{i=1}^m$, and refer to it as PPM(β) for ease of notations.

3. **Creator utility.** Creators are reward-seeking individuals who try to maximize the expected traffic for their created content while carefully balancing the costs. Under PPM(β) and when each creator i produces x_i copies of content, the platform will allocate to i the amount of traffic from u_j proportional to $\mathbb{E}_{\epsilon_{ij} \sim \mathcal{N}(0, \sigma^2)} [x_i e^{\beta_j(w_{ij} + \epsilon_{ij})}] = x_i e^{\beta_j w_{ij}} \cdot e^{\frac{\sigma^2 \beta_j^2}{2}}$. Here we should note that x_i does not suggest the creator would create x_i pieces of identical content, but amount of content following his/her expertise. Therefore, we formulate creator i 's utility function as the following:

$$\begin{aligned} u_i(x_i, \mathbf{x}_{-i}; \beta) &= \sum_{j=1}^m \left(\frac{\mathbb{E}[x_i e^{\beta_j \sigma_{ij}}]}{\sum_{k=1}^n \mathbb{E}[x_k e^{\beta_j \sigma_{kj}}]} \right) - c_i(x_i) \\ &= \sum_{j=1}^m \left(\frac{x_i e^{\beta_j w_{ij}}}{\sum_{k=1}^n x_k e^{\beta_j w_{kj}}} \right) - c_i(x_i), \end{aligned} \quad (1)$$

where $\mathbf{x}_{-i} \in \mathbb{R}_{\geq 0}^{n-1}$ denotes the strategy profile of all creators except i .

The C^4 game models the scenario where creators are aware of their expertise (i.e., what topic to create), and compete purely on *creation quantity*. This concept is akin to the extensively studied Cournot competition [14] model in economics, where firms independently determine the output of homogeneous products at different costs. However, our model diverges from the classic Cournot competition in several key aspects. First of all, they have different revenue functions: in Cournot competition, the revenue for each firm is calculated as the product of price and production quantity, and the price only depends on all firms' joint strategy. In contrast, in the C^4 game, the gain of a creator not only relies on other creators' decision but also hinges on the platform's traffic allocation algorithm. In addition, while Cournot competition typically incorporates only linear cost functions, the C^4 game accommodates general convex cost functions, offering a more nuanced reflection of the real cost faced by creators. These distinctions highlight the unique aspects of our model, while maintaining a conceptual link to traditional economic theories of competition.

Research questions: Under the C^4 framework, an important and natural research question is how we can predict creators' strategic choices in the competition. This is a fundamental question in game theory and we employ the concept of Pure Nash Equilibrium (PNE) [40] to characterize the outcome of C^4 games. The definition of PNE is given by the following:

Definition 1. A joint strategy profile of all creators $\mathbf{x}^* = (x_1^*, \dots, x_n^*)$ forms a pure Nash equilibrium (PNE), if for every creator i , x_i^* is a best response strategy that maximizes u_i given other creators' strategy \mathbf{x}_{-i} ; formally,

$$u_i(x_i^*, \mathbf{x}_{-i}^*) \geq u_i(x_i, \mathbf{x}_{-i}^*) \text{ for every } x_i \in \mathbb{R}_{\geq 0}, \forall i \in [n]. \quad (2)$$

In other word, a PNE represents a stable state when everyone is satisfied with their strategies and does not want to deviate. As we will demonstrate in the subsequent section, a PNE always exists in C^4 under any PPM(β) and can be computed efficiently. This finding forms the basis for our further theoretical analysis and empirical simulations.

In addition to the predictability and stability of creators' production strategies, it is equally crucial for platform designers to develop metrics that encourage the prosperity of the content ecosystem. They must also devise algorithmic solutions to optimize these metrics, balancing the trade-off between engagement of users and creators, using the available "knob" PPM(β). We will explore these issues in the upcoming technical discussions.

4 The PNEs of C^4 Games and Their Properties

As widely known, the PNE does not always exist [18, 23, 26]. However, our first main result establishes that, under mild assumptions, C^4 always admits a unique PNE.

⁴The randomness in matching results may stem from either the imperfect estimation of the preference score w_{ij} or intentionally injected exploration strength based on the intervention mechanism. In this paper, we focus on the latter source of randomness and analyze how such exploration strength might affect the outcomes.

Theorem 1. For any C^4 instance $\mathcal{G}(n, m, M, \{c_i\}_{i=1}^n, \beta)$. If each c_i is convex in x_i , \mathcal{G} admits a unique PNE.

Note that the primary challenge in proving Theorem 1 is to show \mathcal{G} is a strictly monotone game, whereas the existence and uniqueness of PNE in such games is a classic result from [44]. Theorem 1 is interesting from multiple perspectives. First, it strictly generalizes previous equilibrium existence results in classic Tullock contest [42, 13], which corresponds to the special case when $m = 1$ and $w_{11} = \dots = w_{n1}$. Second, the fact that C^4 games are monotone is significant because it is well-known that the PNE of strictly monotone games can be found efficiently. For example, many natural multi-agent online learning dynamics such as mirror descent [7], accelerated optimistic gradient [8], and payoff-based learning [47] guarantee the last-iterate convergence to the unique PNE in strictly monotone games, even when players have mere zeroth order feedback about their utility functions. These results suggest that the PNE of \mathcal{G} is achievable if all creators use a reasonable update rule in their strategies. This observation not only makes this equilibrium a plausible prediction of real-world competition but also paves the way to our simulation-based studies in our experiments, where we use multi-agent mirror descent with perfect gradient to numerically solve the PNE of C^4 .

In addition to the existence and uniqueness properties, the following corollary characterizes the first-order characterization of \mathcal{G} 's PNE.

Corollary 1. The unique PNE $\mathbf{x}^*(\beta)$ of any $\mathcal{G}(n, m, M, \{c_i\}_{i=1}^n, \beta)$ satisfies the following first-order condition:

$$\left. \frac{\partial u_i}{\partial x_i} \right|_{\mathbf{x}=\mathbf{x}^*(\beta)} = 0, \quad 1 \leq i \leq n. \quad (3)$$

Since we have already shown the existence and uniqueness of \mathcal{G} 's PNE, Corollary 1 follows immediately according to Definition 1. Corollary 1 is useful for establishing further properties of C^4 games in Section 5.

5 The Trade-Off Between User and Creator Engagement

We have established that a unique PNE exists in any C^4 game under any PPM(β) and can be naturally achieved by competing content creators. This raises a crucial question for platform designers: how should the quality of the PNE be evaluated? For any mature UGC platform, it is essential to balance user satisfaction, which is key to short-term prosperity, with creator engagement, which is crucial for long-term sustainability. Within our C^4 framework, this requires the platform designer to generate matching results that not only guarantee high user satisfaction (by improving the average matching quality at PNE) but also stimulate substantial content creation volume (by encouraging creators to increase their production frequency x^* at PNE). We define these two objectives as follows:

Definition 2. For any C^4 instance $\mathcal{G}(n, m, M, \{c_i\}_{i=1}^n, \beta)$, let $\mathbf{x}^*(\beta) = (x_1^*, \dots, x_n^*)$ be the unique PNE under PPM(β). Then the (short-term) total user satisfaction is defined as

$$U(\mathbf{x}^*(\beta); \beta) = \sum_{j=1}^m \sum_{i=1}^n \left(\frac{w_{ij} x_i^* e^{\beta_j w_{ij}}}{\sum_{k=1}^n x_k^* e^{\beta_j w_{kj}}} \right), \quad (4)$$

and the (long-term) total content creation volume is defined as

$$V(\mathbf{x}^*(\beta)) = \sum_{i=1}^n x_i^*. \quad (5)$$

The social welfare of the whole system is measured by a linear combination of U and V , defined as

$$W_\lambda(\mathbf{x}^*(\beta); \beta) = U(\mathbf{x}^*(\beta); \beta) + \lambda V(\mathbf{x}^*(\beta); \beta). \quad (6)$$

If we denote $\pi_j^*(\mathbf{x}^*(\beta), \beta) = \sum_{i=1}^n \left(\frac{w_{ij} x_i^* e^{\beta_j w_{ij}}}{\sum_{k=1}^n x_k^* e^{\beta_j w_{kj}}} \right)$ as the indicator of an individual user's satisfaction or utility, which is the expected matching scores of user j at the PNE under PPM(β). And the total user satisfaction measure $U = \sum_{j=1}^m \pi_j^*$ is the accumulated user utility. We argue that U primarily serves as a metric for short-term welfare evaluation, since it focuses solely on user satisfaction at a specific instance of matching outcomes but does not capture the dynamics of user

engagement over time. This overlooks the crucial fact that sustained user engagement on a platform requires a continuous supply of relevant content, as users can hardly be satisfied by their previously consumed material. This limitation is also evident in U 's mathematical formulation: its value remains unchanged with a rescaling of $\{x_i^*\}$, indicating that it fails to reflect changes in content volume or frequency that might affect long-term user engagement. On the other hand, the long-term prosperity of a UGC platform is fundamentally linked to the engagement of content creators. Therefore, we introduce the total content creation volume V as an indicator of the long-term welfare of the platform.

For the platform designers, it is essential to develop metrics that balance both short-term and long-term considerations. Thus, we propose a hybrid social welfare metric, W , which combines U and V to reflect both user satisfaction and content supply sustainability. However, understanding the mechanisms to optimize U and V independently is critical. In the following sections, we will explore the optimal matching mechanisms tailored to the exclusive objectives of U (short-term) and V (long-term), and then present an efficient algorithm designed to optimize W .

Interestingly, our findings suggest that both U and V exhibit monotonicity with respect to the parameter β , even in scenarios involving a homogeneous user population. This uniform behavior of U and V offers valuable insights into how the adjustment of exploration strength could potentially impact platform performance. Our forthcoming theorem formally characterizes these observations.

Theorem 2. *Consider any C^4 game with $m = 1$. If the elements of $M = [w_1, \dots, w_n]^\top$ are not identical, it holds that:*

1. $U(\beta)$ defined in Eq. (4) is strictly increasing in β .
2. $V(\beta)$ defined in Eq. (5) is strictly decreasing in $\beta \in [\beta_0, +\infty)$ for some $\beta_0 > 0$.

Theorem 2 conveys two significant insights. The first one, though perhaps unsurprising, reveals that improving matching accuracy corresponds to an increase in expected user satisfaction. Despite its intuitiveness, this is a strong observation because it holds without relying on any specific structural assumptions about creator cost functions. This means that regardless of the potential complexity in equilibrium structures due to creator costs, and even when the order of x_i^* does not align with a creator's capability w_i , the metric U is still monotonically increasing with respect to β . The second insight may surprise some readers: it suggests that while keep increasing the matching accuracy motivates some creators to produce more content, it demotivates others, resulting in a net decrease in the overall volume of content creation. This finding illustrates an intrinsic trade-off between short-term matching accuracy and long-term content supply: strategies that enhance short-term user satisfaction can inadvertently reduce content creation frequency across creators. To the best of our knowledge, this result is novel and has not been discussed in similar studies.

Here is an intuitive explanation for why a large β diminishes creators' willingness to produce content. As the traffic allocation becomes more deterministic, the marginal gain from increasing production frequency diminishes because the amount of traffic accrued is largely determined by the relevance score, rather than volume. In the extreme case where $\beta \rightarrow +\infty$, only the most relevant creator captures all the user traffic, regardless of her production volume. Consequently, due to the presence of production costs, this creator, and others, will only sustain the minimum viable productivity. Conversely, in the other extreme scenario where $\beta = 0$, i.e., user traffic is distributed uniformly among creators irrespective of relevance, the gain for each creator depends solely on their production frequency, prompting a productivity arms race. Clearly, both extremes are suboptimal, but they effectively illustrate the rationale behind our theoretical findings.

Although in Theorem 2 we consider the game instance \mathcal{G} with $m = 1$ as a representative snapshot of how creators compete for a single unit of user traffic (i.e., homogeneous user population), extending the time frame to encompass a sequence of heterogeneous users suggests that the observed trade-off between U and V remains consistent. In our experiments, we will demonstrate this trade-off in broader settings through simulations, e.g. when $m > 1$ and with various complex user distributions.

The proof of Theorem 2, while delivering a clear message, is far from trivial. Since the dependencies of U and V on β are indirectly linked through x^* , which lacks a closed form, the derivation of their derivatives with respect to β necessitates the use of the implicit function theorem [37] to articulate the derivative of x^* with respect to β . This involves a complex matrix inverse, which we simplify using the Sherman–Morrison formula [46] due to its structure being a diagonal matrix with a rank-one update. This proof technique not only supports our theorem but also inspires a novel first-order

optimization approach to address the hybrid social welfare optimization discussed in Section 7. Detailed proofs are provided in Appendix A.2.

6 Finding the Optimal Trade-off through Optimization

Our theory thus far indicates that optimizing both user satisfaction U and creator engagement V is non-trivial, even when the user population is homogeneous, as achieving the optimal of U and V simultaneously is impossible. Consequently, an essential and intriguing question arises within any specific competitive environment C^4 : how can we identify the optimal trade-off between these two factors by optimizing any given welfare metric W_λ ?

Generally, the welfare metric W is influenced by three factors: the platform's algorithmic recommendation policy $\text{PPM}(\beta)$, the resulting content creation profile \mathbf{x}^* at the PNE induced by β , and the relevance matrix M . Thus, we can formulate the resulting optimization problem (OP) as follows:

$$\begin{aligned} \text{Find} \quad & \arg \max_{\beta \in \mathbb{R}_{\geq 0}^m} W_\lambda(\mathbf{x}^*(\beta), \beta) \\ \text{s.t.} \quad & \mathbf{x}^*(\beta) \text{ is the PNE of } \mathcal{G}. \end{aligned} \quad (7)$$

In general, OP (7) presents a formidable challenge, as solving for a PNE of a game is known to be difficult [15]. Fortunately, the nice structure of C^4 allows us to utilize the implicit characterization of the PNE detailed in Corollary 1 to tackle OP (7) effectively. In the subsequent section, we demonstrate that the gradient of W_λ w.r.t. β can be explicitly computed.

6.1 The Derivation of Exact Gradient

According to the chain rule, the first-order gradient of W_λ w.r.t. β can be expressed as

$$\frac{dW_\lambda}{d\beta} = \frac{dU(\mathbf{x}^*(\beta), \beta)}{d\beta} + \lambda \frac{dV(\mathbf{x}^*(\beta))}{d\beta} = \left(\frac{\partial U}{\partial \mathbf{x}^*} + \lambda \frac{\partial V}{\partial \mathbf{x}^*} \right) \cdot \frac{d\mathbf{x}^*}{d\beta} + \frac{\partial U}{\partial \beta}. \quad (8)$$

The evaluation of the gradient of W_λ relies on the calculation of three vectors, $\frac{\partial U}{\partial \mathbf{x}^*}, \frac{\partial V}{\partial \mathbf{x}^*} \in \mathbb{R}^{1 \times n}$, and $\frac{\partial U}{\partial \beta} \in \mathbb{R}^{1 \times m}$, as well as a Jacobian matrix $\frac{d\mathbf{x}^*}{d\beta} \in \mathbb{R}^{n \times m}$. The computations of $\frac{\partial U}{\partial \mathbf{x}^*}, \frac{\partial V}{\partial \mathbf{x}^*}$ and $\frac{\partial U}{\partial \beta}$ are straightforward and computationally light, which position the main challenge as the computation of $\frac{d\mathbf{x}^*}{d\beta}$. Fortunately, the first-order characterization of \mathbf{x}^* by Corollary 1 enables us to express the gradient of \mathbf{x}^* w.r.t. β using implicit function derivation [37] as $\frac{d\mathbf{x}^*}{d\beta} = -\left(\frac{\partial F}{\partial \mathbf{x}^*}\right)^{-1} \cdot \frac{\partial F}{\partial \beta}$, where $F(\mathbf{x}, \beta) = \left(\frac{\partial u_i}{\partial x_i}\right)_{i=1}^n$ is an n -valued function, and both $\frac{\partial F}{\partial \mathbf{x}^*}$ and $\frac{\partial F}{\partial \beta}$ are matrices of dimensions $n \times n$ and $n \times m$, respectively. The following proposition provides the exact formula for the gradient. The calculation is straightforward and we omit the detailed derivation.

Proposition 1. *Let $\mathbf{x}^* = (x_1^*, \dots, x_n^*)$ be the PNE of $\mathcal{G}(n, m, M, \{c_i\}_{i=1}^n, \beta)$. Then, the Jacobian matrix of \mathbf{x}^* as a function of β is*

$$\frac{d\mathbf{x}^*}{d\beta} = (D + YZ^\top)^{-1} B \in \mathbb{R}^{n \times m}, \quad (9)$$

where D is an $n \times n$ diagonal matrix given by

$$D = \text{diag} \left(c_1'' + \sum_{j=1}^m \frac{P_{1j}^2}{x_1^{*2}}, \dots, c_n'' + \sum_{j=1}^m \frac{P_{nj}^2}{x_n^{*2}} \right), \quad (10)$$

and B, Y, Z are $\mathbb{R}^{n \times m}$ matrices calculated as follows ($1 \leq i \leq n, 1 \leq j \leq m$):

$$Y = \left[\frac{P_{ij}(1 - 2P_{ij})}{x_i^*} \right]_{ij}, \quad Z = \left[\frac{P_{ij}}{x_i^*} \right]_{ij}, \quad B = \left[\frac{P_{ij}(1 - 2P_{ij})}{x_i^*} \cdot \left(w_{ij} - \sum_{k=1}^n w_{kj} P_{kj} \right) \right]_{ij}, \quad (11)$$

where c_i'' is the second-order derivative of creator i 's cost function, and $P_{ij} = \frac{x_i^* \exp(\beta_j w_{ij})}{\sum_{k=1}^n x_k^* \exp(\beta_j w_{kj})}$ is the probability that creator i is matched with user j at the PNE under $\text{PPM}(\beta)$.

6.2 Optimization with Approximated Gradients

Proposition 1 together with Eq. (8) offers us a possibility to directly apply gradient-based approaches for solving OP (7). However, the gradient computation requires the inversion of an $n \times n$ matrix, whose time complexity is $O(n^3)$ and thus too cumbersome. To reduce the computational burden, we propose to approximately compute the gradient using the Sherman–Morrison–Woodbury formula [46] to approximate the matrix inverse, inspired by the specific structure of the RHS of Eq. (9). According to Sherman–Morrison–Woodbury formula, it holds that

$$(D + YZ^\top)^{-1} = D^{-1} - D^{-1}Y(I + Z^\top D^{-1}Y)^{-1}Z^\top D^{-1}, \quad (12)$$

and the computation of the RHS of Eq. (12) now requires a time complexity of $O(n^2m + nm^2 + m^3)$. However, the size of the user population m in practical scenarios is often even larger than n . To efficiently compute the RHS of Eq. (12), we propose a method to “sketch” the matrices Y and Z by sampling a subset of users. Initially, each column of Y and Z corresponds to a user index j . We begin by sampling a sub-population of \mathcal{X} , indexed by \mathcal{I} , with $|\mathcal{I}| = \tilde{m} = \lceil \delta m \rceil$, where $\delta \in (0, 1]$ denotes the sampling rate. With this sampled index set \mathcal{I} , we construct matrices $\tilde{Y}, \tilde{Z} \in \mathbb{R}^{n \times \tilde{m}}$, where the (i, j) -th entries are defined as follows:

$$\tilde{Y}_{ij} = \frac{P_{ij'}(1 - 2P_{ij'})}{x_i^*}, \quad \tilde{Z}_{ij} = \frac{P_{ij'}}{x_i^*}, \quad (13)$$

with j' being uniformly sampled from \mathcal{I} . Given that \tilde{Y}, \tilde{Z} now possess reduced ranks of $\lceil \delta m \rceil$, the computational complexity of evaluating $(D + \tilde{Y}\tilde{Z}^\top)^{-1}$ is significantly lowered to $O(n^2\tilde{m} + n\tilde{m}^2 + \tilde{m}^3)$. Algorithm 1 describes the steps for addressing OP (7).

Algorithm 1: Approximated Gradient Descent for Solving OP (7).

Input: The environment specified by \mathcal{G} , maximum iteration number T , sample rate δ , learning rate η , initial mechanism $\text{PPM}(\beta)$.

- 1 **for** $t \in [T]$ **do**
- 2 Find the PNE x^* of \mathcal{G} under $\text{PPM}(\beta)$ using Algorithm 2,
- 3 Uniformly sample $\lceil \delta m \rceil$ users from \mathcal{X} and use them to compute matrices \tilde{Y}, \tilde{Z} in Eq. (13),
- 4 Compute the approximated gradient $\frac{dW_\lambda}{d\beta}$ using (8),(9),(12) with sketched matrices \tilde{Y}, \tilde{Z} ,
- 5 Update $\beta = \beta + \eta \frac{dW_\lambda}{d\beta}$.

Algorithm 1 requires solving for the PNE of \mathcal{G} each time when β is updated. To accomplish this, we employ the multi-agent mirror descent method, as proposed in [7] and detailed in Algorithm 2 in Appendix, to serve as a subroutine⁵. To accelerate the convergence of Algorithm 1, the PNE strategy x^* obtained under the previous β is used as the initial strategy for computing the new PNE after updating β . Further implementation details are provided in the experiment section.

7 Experiments

To validate our theoretical findings and demonstrate the performance of Algorithm 1, we conduct simulations on instances of \mathcal{G} constructed from both synthetic data and the MovieLens-1m dataset [28]. In our experiments, Algorithm 2 is employed to solve the PNE for each instance of \mathcal{G} . Below, we first outline the specifications of these two simulation environments and then present our results.

Synthetic environment For the synthetic environment, we construct the user population \mathcal{X} by setting an embedding dimension $d = 32$ and independently sampling 50 cluster centers, denoted as $\{c_1, \dots, c_{50}\}$, from the unit sphere \mathbb{S}^{d-1} . For each center c_i , users belonging to cluster- i are generated by sampling independently from a Gaussian distribution $\mathcal{N}(c_i, 0.5^2 \mathbf{I}_d)$. The sizes of the 50 user clusters are determined uniformly at random, ensuring the total size of \mathcal{X} is $m = 1000$. Similarly, $n = 200$ creators are generated, and the relevance matrix $M \in \mathbb{R}^{n \times m}$ is defined by the dot product between each user-creator pair, which are then normalized to the range $[0, 1]$. This synthetic

⁵In [7], the algorithm is guaranteed to converge to PNE under zeroth order feedback. Here we use the perfect gradient as input and thus the convergence is also guaranteed.




Figure 1: The left and the middle panel: the empirical distributions of content creation frequency x_i^* and each user’s individual utility π_j^* . Different colors represent results for PNEs induced by different β . Right: the total content creation V and total user satisfaction U obtained under different β . Error bars obtained from 10 independently generated environments.

dataset encapsulates a class of clustered user and creator preference distributions. On the creators’ side, their cost functions are set to $c_i(x) = c_i x^\rho$, with the default $\rho = 1.5$. The marginal costs $\{c_i\}_{i=1}^n$ are randomly sampled from a uniform distribution $\mathcal{U}[0.1, 0.5]$.

Environment constructed from MovieLens-1m dataset We use deep matrix factorization [22] to train user and movie embeddings (with dimension set to 32) by fitting the observed ratings in the range of 1 to 5. To ensure the quality of the trained embeddings, we performed a 5-fold cross-validation and obtained an averaged RMSE=0.739 on the test sets. With the same hyper-parameter, we train the user/item embeddings with the complete dataset. We randomly select $m = 1000$ user embeddings to construct the population \mathcal{X} and $n = 200$ movie embeddings as the creator profiles. Similarly, The relevance matrix $M \in \mathbb{R}^{n \times m}$ is given by the dot product between each user-creator pair normalized to $[0, 1]$ and creators’ cost functions are the same as we specified in the synthetic environment.

7.1 The Empirical Trade-Offs Between U and V

Figure 1 illustrates the content creation frequency x_i^* , user utility π_j^* , and their corresponding aggregated values $U = \sum_j \pi_j^*$, $V = \sum_i x_i^*$ under the PNE induced by different homogeneous β (i.e., all users share the same β). The result in the right panel shows that a larger β enhances overall user satisfaction U but undermines total content creation V . As β increases, the drop in V becomes more significant. This empirical finding supports Theorem 2 and suggests it holds under broader settings without the assumptions on creator cost function and user population structure. The left and middle plots illustrate each creator i ’s creation frequency x_i^* and each user j ’s utility π_j^* at the PNE, such that both x_i^* and π_j^* are rearranged in descending order. They show that when β is shared across all users, its change affects x_i^* and π_j^* in the same direction.

7.2 The Optimal PPM(β) Found by Algorithm 1

Next, we use Algorithm 1 to find the optimal PPM(β) and investigate the properties of the optimal β . We set $\lambda = 0.5$ and aim to maximize the objective $W_\lambda = U + 0.5V$, more results under different choices of λ can be found in Appendix B. The initial β is set to $(100, \dots, 100)$, representing a nearly deterministic matching for every user. Algorithm 1 is then run to update β . The sample rate and learning rate are set to $\delta = 0.1, \eta = 200$. In addition to searching for personalized β_j for each user j , we also attempt to find a homogeneous β (i.e., a fixed β_j for each j) using Algorithm 1⁶.

The first and third panels in Figure 2 show the evolution of W_λ during the optimization process in both synthetic and MovieLens environments. As illustrated, Algorithm 1 successfully finds a better PPM(β) compared to the baseline of exact matching for all users, with a significant gain of over 20% in the welfare metric. Furthermore, in both environments, personalized β leads to a slightly better outcome compared to homogeneous β .

The second and fourth panels depict the optimal β_j for each user j , arranged in descending order. These panels provide insights into how such a mechanism achieves better trade-offs. For each user

⁶The gradient of W_λ with respect to a homogeneous β can be readily obtained by summing all the partial derivatives of W_λ with respect to β_j .

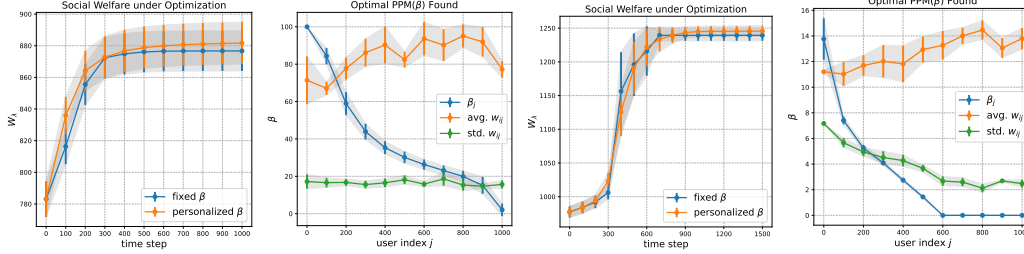


Figure 2: Panel 1,2: social welfare improving curve under Algorithm 1, and the distribution of the obtained optimal β_j in the synthetic environment. Panel 3,4: the same plots in the MovieLens environment. $\lambda = 0.5$.

index j on the x -axis, we also plot the average and the standard deviation of the relevance scores $\{w_{ij}\}_{i=1}^n$ associated with each user j over all creators, shown as orange and green lines. Based on the definition, users with smaller average scores and higher standard deviations are considered more “picky” or selective, indicating high relevance scores with a small group of creators and low scores with many others. Conversely, users with higher average scores and smaller standard deviations are less selective and more open to exploration. The results show that the optimal β tends to increase the exploration strengths (by deploying smaller β_j) for less selective users. This approach is intuitive, as it safely increases exploration while minimizing losses in user engagement.

8 Conclusion

In this work, we introduced a new game-theoretical model C^4 (Cournot Content Creation Competition) to explore how creators strategically determine their creation frequency under a UGC platform’s recommendation algorithm. Our investigations reveal a critical balance between user satisfaction and creator engagement, mediated by the exploration strength of the recommendation. The existence and uniqueness of the PNE of C^4 games provide a predictive framework for assessing the effects of algorithmic choices on content diversity and volume. Through both theoretical analysis and empirical simulations, we demonstrated how varying the exploration strength can either enhance user engagement at the cost of reduced content diversity or encourage richer content creation at the expense of immediate user satisfaction. These findings disclose the delicate trade-offs platform designers face and highlight the utility of our model as a pre-deployment audit tool for optimizing recommendation algorithms to balance platforms’ long-term and short-term objectives.

While our C^4 model offers insights into strategic differentiation among creators regarding production quantity, it relies on a simplified assumption that creators maintain a fixed niche, consistently producing content on the same topic with similar quality. This assumption, though useful for modeling purposes, may be restrictive in real-world scenarios where creators dynamically adjust topics, vary content quality, and scale production quantity. Exploring the dynamics where creators compete across heterogeneous dimensions—such as topic variety, content quality, and production quantity—would be a valuable direction for future research. We leave this intriguing problem for future work.

Acknowledgment. This work is supported in part by the NSF Award IIS-2128019, NSF Award CCF-2303372, AI2050 program at Schmidt Sciences (Grant G-24-66104) and Army Research Office Award W911NF-23-1-0030.

References

- [1] Arpit Agarwal and William Brown. Online recommendations for agents with discounted adaptive preferences. *arXiv preprint arXiv:2302.06014*, 2023.
- [2] Omer Ben-Porat and Moshe Tennenholtz. Shapley facility location games. In *International Conference on Web and Internet Economics*, pages 58–73. Springer, 2017.

- [3] Omer Ben-Porat and Moshe Tennenholtz. A game-theoretic approach to recommendation systems with strategic content providers. *Advances in Neural Information Processing Systems*, 31, 2018.
- [4] Erdem Biyik, Fan Yao, Yinlam Chow, Alex Haig, Chih-wei Hsu, Mohammad Ghavamzadeh, and Craig Boutilier. Preference elicitation with soft attributes in interactive recommendation. *arXiv preprint arXiv:2311.02085*, 2023.
- [5] Jesús Bobadilla, Fernando Ortega, Antonio Hernando, and Abraham Gutiérrez. Recommender systems survey. *Knowledge-based systems*, 46:109–132, 2013.
- [6] Kianté Brantley, Zhichong Fang, Sarah Dean, and Thorsten Joachims. Ranking with long-term constraints. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*, pages 47–56, 2024.
- [7] Mario Bravo, David Leslie, and Panayotis Mertikopoulos. Bandit learning in concave n-person games. *Advances in Neural Information Processing Systems*, 31, 2018.
- [8] Yang Cai and Weiqiang Zheng. Doubly optimal no-regret learning in monotone games. In *International Conference on Machine Learning*, pages 3507–3524. PMLR, 2023.
- [9] Ioannis Caragiannis and Alexandros A Voudouris. Welfare guarantees for proportional allocations. *Theory of Computing Systems*, 59:581–599, 2016.
- [10] Micah Carroll, Dylan Hadfield-Menell, Stuart Russell, and Anca Dragan. Estimating and penalizing preference shift in recommender systems. In *Proceedings of the 15th ACM Conference on Recommender Systems*, pages 661–667, 2021.
- [11] Sarah H Cen, Andrew Ilyas, Jennifer Allen, Hannah Li, and Aleksander Madry. Measuring strategization in recommendation: Users adapt their behavior to shape future content. *arXiv preprint arXiv:2405.05596*, 2024.
- [12] Sarah H Cen, Andrew Ilyas, and Aleksander Madry. User strategization and trustworthy algorithms. *arXiv preprint arXiv:2312.17666*, 2023.
- [13] Richard Cornes and Roger Hartley. Asymmetric contests with general technologies. *Economic theory*, 26:923–946, 2005.
- [14] Antoine Augustin Cournot. *Recherches sur les principes mathématiques de la théorie des richesses*, volume 48. L. Hachette, 1838.
- [15] Constantinos Daskalakis, Paul W Goldberg, and Christos H Papadimitriou. The complexity of computing a nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.
- [16] Sarah Dean, Evan Dong, Meena Jagadeesan, and Liu Leqi. Recommender systems as dynamical systems: Interactions with viewers and creators. In *Workshop on Recommendation Ecosystems: Modeling, Optimization and Incentive Design*, 2024.
- [17] Sarah Dean and Jamie Morgenstern. Preference dynamics under personalized recommendations. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 795–816, 2022.
- [18] Gerard Debreu. A social equilibrium existence theorem. *Proceedings of the National Academy of Sciences*, 38(10):886–893, 1952.
- [19] Itay Eilat and Nir Rosenfeld. Performative recommendation: diversifying content via strategic incentives. In *International Conference on Machine Learning*, pages 9082–9103. PMLR, 2023.
- [20] Christian Ewerhart. Mixed equilibria in tullock contests. *Economic Theory*, 60:59–71, 2015.
- [21] Christian Ewerhart. The lottery contest is a best-response potential game. *Economics Letters*, 155:168–171, 2017.
- [22] Jicong Fan and Jieyu Cheng. Matrix completion by deep matrix factorization. *Neural Networks*, 98:34–41, 2018.

- [23] Ky Fan. Fixed-point and minimax theorems in locally convex topological linear spaces. *Proceedings of the National Academy of Sciences*, 38(2):121–126, 1952.
- [24] Arpita Ghosh and Patrick Hummel. Learning and incentives in user-generated content: Multi-armed bandits with endogenous arms. In *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*, pages 233–246, 2013.
- [25] Arpita Ghosh and Preston McAfee. Incentivizing high-quality user-generated content. In *Proceedings of the 20th international conference on World wide web*, pages 137–146, 2011.
- [26] Irving L Glicksberg. A further generalization of the kakutani fixed theorem, with application to nash equilibrium points. *Proceedings of the American Mathematical Society*, 3(1):170–174, 1952.
- [27] Angela Glotfelter. Algorithmic circulation: how content creators navigate the effects of algorithms on their work. *Computers and composition*, 54:102521, 2019.
- [28] F Maxwell Harper and Joseph A Konstan. The movielens datasets: History and context. *Acm transactions on interactive intelligent systems (tiis)*, 5(4):1–19, 2015.
- [29] Thomas Hodgson. Spotify and the democratisation of music. *Popular Music*, 40(1):1–17, 2021.
- [30] Jiri Hron, Karl Krauth, Michael I Jordan, Niki Kilbertus, and Sarah Dean. Modeling content creator incentives on algorithm-curated platforms. *arXiv preprint arXiv:2206.13102*, 2022.
- [31] Qinlu Hu, Ni Huang, and Renyu Philip Zhang. Viewer traffic allocation for small creator development: Experimental evidence from short-video platforms. *Available at SSRN 4888995*, 2024.
- [32] Xinyan Hu, Meena Jagadeesan, Michael I Jordan, and Jacob Steinhard. Incentivizing high-quality content in online recommender systems. *arXiv preprint arXiv:2306.07479*, 2023.
- [33] Nicole Immorlica, Meena Jagadeesan, and Brendan Lucier. Clickbait vs. quality: How engagement-based optimization shapes the content landscape in online platforms. In *Proceedings of the ACM on Web Conference 2024*, pages 36–45, 2024.
- [34] Meena Jagadeesan, Nikhil Garg, and Jacob Steinhardt. Supply-side equilibria in recommender systems. *Advances in Neural Information Processing Systems*, 36, 2024.
- [35] Dimitris Kalimeris, Smriti Bhagat, Shankar Kalyanaraman, and Udi Weinsberg. Preference amplification in recommender systems. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 805–815, 2021.
- [36] Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37, 2009.
- [37] Steven George Krantz and Harold R Parks. *The implicit function theorem: history, theory, and applications*. Springer Science & Business Media, 2002.
- [38] Tao Lin, Kun Jin, Andrew Estornell, Xiaoying Zhang, Yiling Chen, and Yang Liu. User-creator feature dynamics in recommender systems with dual influence. *arXiv preprint arXiv:2407.14094*, 2024.
- [39] Martin Mladenov, Elliot Creager, Omer Ben-Porat, Kevin Swersky, Richard Zemel, and Craig Boutilier. Optimizing long-term social welfare in recommender systems: A constrained matching approach. In *International Conference on Machine Learning*, pages 6987–6998. PMLR, 2020.
- [40] John F Nash Jr. Equilibrium points in n-person games. *Proceedings of the national academy of sciences*, 36(1):48–49, 1950.
- [41] Thành Nguyen and Milan Vojnovic. Weighted proportional allocation. *ACM SIGMETRICS Performance Evaluation Review*, 39(1):133–144, 2011.

- [42] J David Pérez-Castrillo and Thierry Verdier. A general analysis of rent-seeking games. *Public choice*, 73(3):335–350, 1992.
- [43] Siddharth Prasad, Martin Mladenov, and Craig Boutilier. Content prompting: Modeling content provider dynamics to improve user welfare in recommender ecosystems. *arXiv preprint arXiv:2309.00940*, 2023.
- [44] J Ben Rosen. Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica: Journal of the Econometric Society*, pages 520–534, 1965.
- [45] Marcelo Luis Barbosa dos Santos. The “so-called” ugc: an updated definition of user-generated content in the age of social media. *Online Information Review*, 46(1):95–113, 2022.
- [46] Jack Sherman. Adjustment of an inverse matrix corresponding to changes in the elements of a given column or row of the original matrix. *Annu. Math. Statist.*, 20:621, 1949.
- [47] Tatiana Tatarenko and Maryam Kamgarpour. Bandit learning in convex non-strictly monotone games. *arXiv preprint arXiv:2009.04258*, 2020.
- [48] Gordon Tullock. Efficient rent seeking. In *Toward a theory of the rent-seeking society*. College Station, TX: Texas A&M University Press, 1980.
- [49] Renzhe Xu, Haotian Wang, Xingxuan Zhang, Bo Li, and Peng Cui. Ppa-game: Characterizing and learning competitive dynamics among online content creators. *arXiv preprint arXiv:2403.15524*, 2024.
- [50] Fan Yao, Chuanhao Li, Denis Nekipelov, Hongning Wang, and Haifeng Xu. Learning from a learning user for optimal recommendations. In *International Conference on Machine Learning*, pages 25382–25406. PMLR, 2022.
- [51] Fan Yao, Chuanhao Li, Denis Nekipelov, Hongning Wang, and Haifeng Xu. Learning the optimal recommendation from explorative users. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 9457–9465, 2022.
- [52] Fan Yao, Chuanhao Li, Denis Nekipelov, Hongning Wang, and Haifeng Xu. How bad is top- k recommendation under competing content creators? In *International Conference on Machine Learning*. PMLR, 2023.
- [53] Fan Yao, Chuanhao Li, Denis Nekipelov, Hongning Wang, and Haifeng Xu. Human vs. generative ai in content creation competition: Symbiosis or conflict? In *International Conference on Machine Learning*. PMLR, 2024.
- [54] Fan Yao, Chuanhao Li, Karthik Abinav Sankararaman, Yiming Liao, Yan Zhu, Qifan Wang, Hongning Wang, and Haifeng Xu. Rethinking incentives in recommender systems: Are monotone rewards always beneficial? *Advances in Neural Information Processing Systems*, 36, 2024.
- [55] Fan Yao, Yiming Liao, Mingzhe Wu, Chuanhao Li, Yan Zhu, James Yang, Jingzhou Liu, Qifan Wang, Haifeng Xu, and Hongning Wang. User welfare optimization in recommender systems with competing content creators. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 3874–3885, 2024.
- [56] Zhiyu Zeng, Zhiqi Zhang, Dennis Zhang, and Tat Chan. The impact of recommender systems on content consumption and production: Evidence from field experiments and structural modeling. *Available at SSRN*, 2024.
- [57] Banghua Zhu, Sai Praneeth Karimireddy, Jiantao Jiao, and Michael I Jordan. Online learning in a creator economy. *arXiv preprint arXiv:2305.11381*, 2023.
- [58] Wei Zhuang, Qingfeng Zeng, Yu Zhang, Chunmei Liu, and Weiguo Fan. What makes user-generated content more helpful on social media platforms? insights from creator interactivity perspective. *Information processing & management*, 60(2):103201, 2023.

Appendix to *Unveiling User Satisfaction and Creator Productivity Trade-Offs in Recommendation Platforms*

A Omitted Proofs

A.1 Proof of Theorem 1

Proof. First of all, we argue that given any $\mathcal{G}(n, m, M, \{c_i\}, \beta)$, for any creator i , there exists an $\delta_i > 0$ such that any $x_i \in [0, \delta_i]$ cannot be an equilibrium strategy. This is because given any $\mathbf{x}_{-i} \in \mathbb{R}_{\geq 0}^{n-1}$, u_i as a function of x_i has a continuous and strictly positive gradient at $x_i = 0$, meaning

that there exists a $\delta_i > 0$ such that $\left. \frac{\partial u_i}{\partial x_i} \right|_{x_i=t} > 0, \forall t \in [0, \delta_i]$ regardless of what other creators' strategies are. In other word, for any $x_i \leq \delta_i$, creator i can always increase her strategy to strictly improve her utility. As a result, any potential PNE \mathbf{x}^* must satisfy that $x_i^* \geq \delta_i$.

On the other hand, since $c_i(x_i) \rightarrow +\infty$ when $x_i \rightarrow +\infty$ but the traffic gain for each creator is at most m , we have $u_i(x_i, \mathbf{x}_{-i}) \rightarrow -\infty, \forall \mathbf{x}_{-i} \in \mathbb{R}_{\geq 0}^{n-1}$ when $x_i \rightarrow +\infty$. As a result, any equilibrium strategy must also be upper bounded by a uniform constant $\Delta > 0$.

To argue the existence and uniqueness of PNE of \mathcal{G} , in the following we may with out loss of generality restrict each creator i 's strategy set to a convex set $[\delta_i, \Delta]$.

For any fixed $\beta = (\beta_1, \dots, \beta_m)$, let $a_{ij} = \exp(\beta_j w_{ij})$ and Eq. (1) can be simplified to

$$u_i(x_i, \mathbf{x}_{-i}; \beta) = \sum_{j=1}^m \left(\frac{x_i a_{ij}}{\sum_{k=1}^n x_k a_{kj}} \right) - c_i(x_i), \quad (14)$$

For simplicity we denote $g_j(\mathbf{x}) = \sum_{k=1}^n x_k a_{kj}$ and u_i can be expressed as $u_i(x_i, \mathbf{x}_{-i}) = \sum_{j=1}^m \frac{x_i a_{ij}}{g_j(\mathbf{x})} - c_i(x_i)$. Our proof starts from a sufficient condition from [44] for a game to be monotone. A game is said to satisfy the *diagonal strict concavity* (DSC) condition if (1) each player has a concave utility function in his own strategy in a convex strategy space; and (2) there exists *some* non-zero parameter $\lambda = (\lambda_1, \dots, \lambda_n)$ such that the Hessian matrix given by

$$H_{kl}(\mathbf{x}; \lambda) \triangleq \frac{\lambda_k}{2} \frac{\partial^2 u_k(\mathbf{x})}{\partial x_k \partial x_l} + \frac{\lambda_l}{2} \frac{\partial^2 u_l(\mathbf{x})}{\partial x_l \partial x_k} \quad (15)$$

is *strictly* negative-definite. In [44], it is shown that any game satisfying λ -DSC condition has a unique pure Nash equilibrium (PNE); such a game is often referred to as monotone games.

First of all, we already argued that each creator i 's strategy set is $[\delta_i, \Delta]$, which is a convex set. Core to our proof is to show that game \mathcal{G} is 1-DSC under the theorem conditions. Direct calculation shows that for any $1 \leq k \leq l \leq n$,

$$\begin{aligned} \frac{\partial^2 u_k(\mathbf{x})}{\partial x_k \partial x_l} &= \frac{a_{kj} a_{lj}}{g_j^3} \cdot (-g_j + 2a_{kj} x_k), \\ \frac{\partial^2 u_l(\mathbf{x})}{\partial x_l \partial x_k} &= \frac{a_{kj} a_{lj}}{g_j^3} \cdot (-g_j + 2a_{lj} x_l), \end{aligned}$$

and therefore the Hessian matrix of \mathcal{G} specified by the RHS of Eq. (15) is equal to

$$\begin{aligned} -[H(\mathbf{x})] &= \sum_{j=1}^m g_j^{-3} \begin{bmatrix} a_{1j} \\ a_{2j} \\ \vdots \end{bmatrix} \begin{bmatrix} 2 \sum_{i \neq 1} x_i a_{ij} & \sum_{i \notin \{1,2\}} x_i a_{ij} & \cdots \\ \sum_{i \notin \{1,2\}} x_i a_{ij} & 2 \sum_{i \neq 2} x_i a_{ij} & \cdots \\ \vdots & \vdots & \ddots \end{bmatrix} [a_{1j}, a_{2j}, \dots] + \begin{bmatrix} \frac{\partial^2 c_1}{\partial x_1^2} & 0 & \cdots \\ 0 & \frac{\partial^2 c_2}{\partial x_2^2} & \cdots \\ \vdots & \vdots & \ddots \end{bmatrix} \\ &\triangleq \sum_{j=1}^m g_j^{-3} \mathbf{a}_j H_j \mathbf{a}_j^\top + H_0, \end{aligned} \quad (16)$$

where g_j in the above expressions denotes $g_j(\mathbf{x})$, vector $\mathbf{a}_j = (a_{1j}, \dots, a_{nj})^\top$. We can see that if all the cost functions are strictly convex, the second diagonal matrix H_0 in the RHS of Eq. (16) is strictly positive-definite (PD). Therefore, it suffices to show that (1) for all $j \in [m]$, H_j is PD. To see this, let $z_i = x_i a_{ij}$ and we show that for any $\mathbf{y} = (y_1, \dots, y_n) \in \mathbb{R}^n$, $\mathbf{y}^\top H_j \mathbf{y} \geq 0$, and the equality holds if and only if $\mathbf{y} = \mathbf{0}$. In fact, note that

$$\begin{aligned}
\mathbf{y}^\top H_j \mathbf{y} &= 2 \sum_{i=1}^n y_i^2 \left(\sum_{j \neq i} z_j \right) + 2 \sum_{i < j} y_i y_j \left(\sum_{k \notin \{i, j\}} z_k \right) \\
&= \sum_{i=1}^n y_i^2 \left(\sum_{j \neq i} z_j \right) + \left[\sum_{i=1}^n y_i^2 \left(\sum_{j \neq i} z_j \right) + 2 \sum_{i < j} y_i y_j \left(\sum_{k \notin \{i, j\}} z_k \right) \right] \\
&= \sum_{i=1}^n y_i^2 \left(\sum_{j \neq i} z_j \right) + \sum_{k=1}^n z_k \left[\sum_{j \neq k} y_j^2 + 2 \sum_{i < j, i \neq k, j \neq k} y_i y_j \right] \\
&= \sum_{i=1}^n y_i^2 \left(\sum_{j \neq i} z_j \right) + \sum_{i=1}^n z_i \left(\sum_{j \neq i} y_j \right)^2 \geq 0. \tag{17}
\end{aligned}$$

Because x_i and a_i are all strictly positive, each z_i must also be strictly positive. Hence, Eq. (17) can take value zero if and only if $y_i = 0, \forall i \in [n]$. Therefore, H_j is PD for any $j \in [m]$, which completes the proof. \square

A.2 Proof of Theorem 2

First let's recall the definition of U, V when $m = 1$:

$$U(\mathbf{x}^*(\beta); \beta) = \sum_{i=1}^n \left(\frac{w_i x_i^* e^{\beta w_i}}{\sum_{k=1}^n x_k^* e^{\beta w_k}} \right), \tag{18}$$

$$V(\mathbf{x}^*) = \sum_{i=1}^n x_i^*. \tag{19}$$

In the following, we prove the monotonicity of $U(\beta)$ and $V(\beta)$ by showing $\frac{d \ln U}{d \beta} > 0$ and $\frac{dV}{d\beta} < 0$, respectively. Before presenting the detailed proof, we first derive some relevant definitions and their properties that will be used in the proof.

For simplicity we omit the superscript $*$ in \mathbf{x}^* and simply use \mathbf{x} to refer to the PNE of \mathcal{G} . When $m = 1$, the creator utility function writes

$$u_i(x_i, x_{-i}) = \frac{x_i e^{\beta w_i}}{\sum_{k=1}^n x_k e^{\beta w_k}} - c_i(x_i), i \in [n]. \tag{20}$$

Let $P_i = \frac{x_i a_i}{\sum_{k=1}^n x_k a_k}$, where $a_i = e^{\beta w_i}$. First of all, we claim that it is without loss of generality to consider the regime where $P_i \leq \frac{1}{3}$. To see this, consider the following two C^4 instances:

$$\begin{aligned}
\mathcal{G}_1(n, m = 1, \mathbf{w} = (w_1, \dots, w_n), \mathbf{c} = (c_1, \dots, c_n), \beta), \\
\mathcal{G}_2(3n, m = 1, [\mathbf{w}, \mathbf{w}, \mathbf{w}], [\mathbf{c}/3, \mathbf{c}/3, \mathbf{c}/3], [\beta, \beta, \beta]).
\end{aligned}$$

Clearly, both games $\mathcal{G}_1, \mathcal{G}_2$ have unique PNE. Let the PNE of \mathcal{G}_1 be denoted by \mathbf{x}_1^* . In \mathcal{G}_2 , since its $3n$ players are divided into three identical groups, its PNE can be represented as $(\mathbf{x}_2^*, \mathbf{x}_2^*, \mathbf{x}_2^*)$, where each group of n players follow the same strategy. Moreover, for any $1 \leq i \leq n$, it is straightforward to observe that the i -th player's utility functions in \mathcal{G}_1 and \mathcal{G}_2 differ only by a multiplicative constant of 3. Consequently, we have $\mathbf{x}_1^* = \mathbf{x}_2^*$. Therefore, for any C^4 instance \mathcal{G}_1 with $m = 1$, we can always

construct an equivalent instance \mathcal{G}_2 that shares the same PNE structure, while ensuring $P_i \leq \frac{1}{3}$. This justifies the assumptions that, without loss of generality, we can take $P_i \leq \frac{1}{3}$.

Another property we need is that the PNE strategy x_i of any player i in \mathcal{G}_1 is bounded in a compact region $[0, L]$ for some constant L , regardless of the values of β . To see this, note that $c_i(x_i)$ is increasing in x_i and goes to infinity as $x_i \rightarrow +\infty$ while P_i is upper bounded by 1. As a result, for any x_i such that $u_i(x_i, \mathbf{x}_{-i}) \leq 1 - c_i(x_i) < -c_i(0) = u_i(0, \mathbf{x}_{-i})$, x_i cannot be a PNE strategy as switching to 0 increases player i 's utility. Therefore, if we take

$$L = \max_{i \in [n]} \{ \inf_x \{ x \geq 0 : 1 - c_i(x_i) < -c_i(0) \} \},$$

it holds that $x_i \leq L, \forall i \in [n]$.

Let $F(\mathbf{x}, \beta) = \left(\frac{\partial u_i}{\partial x_i} \right)_{i=1}^n$ be an n -value function. From Corollary 1 we know \mathbf{x}, β satisfy $F(\mathbf{x}, \beta) = 0$. And Theorem 1 guarantees that for any $\beta \geq 0$, the \mathbf{x} implicitly determined by $F(\mathbf{x}, \beta) = 0$ exists and is unique. Therefore, by the implicit function theorem [37], the derivative of \mathbf{x} w.r.t. β can be written as

$$\frac{d\mathbf{x}}{d\beta} = - \left(\frac{\partial F}{\partial \mathbf{x}} \right)^{-1} \cdot \frac{\partial F}{\partial \beta},$$

where $\left[\frac{\partial F}{\partial \mathbf{x}} \right]_{n \times n}$ is the Jacobian matrix (which is also the Hessian of \mathcal{G} when $\lambda = 1$, see Eq. (15)) and $\frac{\partial F}{\partial \beta} \in \mathbb{R}^{n \times 1}$ is the partial derivative of F w.r.t. β .

The first-order derivative of u_i can be calculated as

$$\frac{\partial u_i}{\partial x_i} = \frac{e^{\beta w_i}}{\sum_{k=1}^n x_k e^{\beta w_k}} - x_i \left(\frac{e^{\beta w_i}}{\sum_{k=1}^n x_k e^{\beta w_k}} \right)^2 - c'_i(x_i), \quad (21)$$

and we can use it to further obtain the following explicit expressions in terms of the derivatives of F :

$$\begin{aligned} \left(\frac{\partial F_i}{\partial x_i} \right) &= -\frac{1}{x_i^2} P_i^2 (1 - 2P_i) - \frac{1}{x_i^2} P_i^2 - c''_i(x_i), \\ \left(\frac{\partial F_i}{\partial x_j} \right) &= -\frac{1}{x_i x_j} P_i P_j (1 - 2P_i), j \neq i, \\ \left(\frac{\partial F_i}{\partial \beta} \right) &= \frac{P_i}{x_i} \cdot (1 - 2P_i) \cdot \left(w_i - \sum_{k=1}^n P_k w_k \right). \end{aligned} \quad (22)$$

Let's define a positive definite diagonal matrix

$$D = \text{diag} \left(\frac{P_1^2}{x_1^2} + c''_1, \dots, \frac{P_n^2}{x_n^2} + c''_n \right).$$

Since $c''_i > 0$, we can introduce variables

$$\delta_i \in (0, 1) \text{ such that } \frac{P_i^2}{x_i^2} + c''_i = \frac{P_i^2}{\delta_i x_i^2}. \quad (23)$$

In addition, let's also define

$$\mathbf{y} = \left(\frac{P_1(1 - 2P_1)}{x_1}, \dots, \frac{P_n(1 - 2P_n)}{x_n} \right)^\top, \mathbf{z} = \left(\frac{P_1}{x_1}, \dots, \frac{P_n}{x_n} \right)^\top, \quad (24)$$

then we have $\frac{\partial F}{\partial \mathbf{x}} = D + \mathbf{y}\mathbf{z}^\top$ and from Sherman–Morrison formula[46], it holds that

$$\begin{aligned} - \left(\frac{\partial F}{\partial \mathbf{x}} \right)^{-1} &= (D + \mathbf{y}\mathbf{z}^\top)^{-1} \\ &= D^{-1} - \frac{D^{-1} \mathbf{y}\mathbf{z}^\top D^{-1}}{1 + \mathbf{z}^\top D^{-1} \mathbf{y}}, \end{aligned} \quad (25)$$

where

$$\begin{aligned} D^{-1} &= \text{diag} \left(\frac{\delta_1 x_1^2}{P_1^2}, \dots, \frac{\delta_n x_n^2}{P_n^2} \right), \\ D^{-1} \mathbf{y} &= \left(\frac{\delta_1 x_1 (1 - 2P_1)}{P_1}, \dots, \frac{\delta_n x_n (1 - 2P_n)}{P_n} \right)^\top, \\ \mathbf{z}^\top D^{-1} &= \left(\frac{\delta_1 x_1}{P_1}, \dots, \frac{\delta_n x_n}{P_n} \right). \end{aligned}$$

Since $P_i \leq \frac{1}{3}$, we have $1 - 2P_i \geq \frac{1}{3} > 0$. With all the notations introduced so far we are now ready to give the formal proof of Theorem 2.

Proof. We prove the monotonicity of $U(\beta)$ and $V(\beta)$ by showing $\frac{d \ln U}{d\beta} > 0$ and $\frac{dV}{d\beta} < 0$.

The monotonicity of $U(\beta)$: The first-order derivative of $\ln U$ w.r.t. β is given by

$$\begin{aligned} \frac{d \ln U}{d\beta} &= \frac{1}{U} \cdot \left(\frac{\partial U}{\partial \mathbf{x}} \cdot \frac{d\mathbf{x}}{d\beta} + \frac{\partial U}{\partial \beta} \right) \\ &= -\frac{1}{U} \cdot \frac{\partial U}{\partial \mathbf{x}} \cdot \left(\frac{\partial F}{\partial \mathbf{x}} \right)^{-1} \cdot \frac{\partial F}{\partial \beta} + \frac{1}{U} \cdot \frac{\partial U}{\partial \beta}. \end{aligned} \quad (26)$$

where $\frac{\partial U}{\partial \mathbf{x}} \in \mathbb{R}^{1 \times n}$ is the partial derivative of U w.r.t. \mathbf{x} . Let $a_i = e^{\beta w_i}$, we will first show $\frac{1}{U} \cdot \frac{\partial U}{\partial \beta} \geq 0$. In fact, calculation shows

$$\begin{aligned} \frac{1}{U} \cdot \frac{\partial U}{\partial \beta} &= \frac{\sum_{k=1}^n x_k a_k}{\sum_{k=1}^n w_k x_k a_k} \cdot \left(\frac{(\sum_{k=1}^n w_k^2 x_k a_k)(\sum_{k=1}^n x_k a_k) - (\sum_{k=1}^n w_k x_k a_k)^2}{(\sum_{k=1}^n x_k a_k)^2} \right) \\ &= \frac{\sum_{k=1}^n w_k^2 x_k a_k}{\sum_{k=1}^n w_k x_k a_k} - \frac{\sum_{k=1}^n w_k x_k a_k}{\sum_{k=1}^n x_k a_k}. \end{aligned} \quad (27)$$

From Cauchy–Schwarz inequality, it holds that

$$\sum_{k=1}^n w_k^2 x_k a_k \cdot \sum_{k=1}^n x_k a_k \geq \left(\sum_{k=1}^n \sqrt{w_k^2 x_k a_k \cdot x_k a_k} \right)^2 = \left(\sum_{k=1}^n w_k x_k a_k \right)^2.$$

Therefore, the RHS of Eq. (27) is greater than or equal to 0. Hence, it suffices to show

$$-\frac{1}{U} \cdot \frac{\partial U}{\partial \mathbf{x}} \cdot \left(\frac{\partial F}{\partial \mathbf{x}} \right)^{-1} \cdot \frac{\partial F}{\partial \beta} > 0. \quad (28)$$

Also note that

$$\begin{aligned} \frac{1}{U} \cdot \frac{\partial U}{\partial x_i} &= \frac{\sum_{k=1}^n x_k a_k}{\sum_{k=1}^n w_k x_k a_k} \cdot \left(\frac{w_i a_i (\sum_{k=1}^n x_k a_k) - a_i \sum_{k=1}^n w_k x_k a_k}{(\sum_{k=1}^n x_k a_k)^2} \right) \\ &= \frac{w_i a_i}{\sum_{k=1}^n w_k x_k a_k} - \frac{a_i}{\sum_{k=1}^n x_k a_k} \\ &= \frac{w_i a_i}{\sum_{k=1}^n w_k x_k a_k} - \frac{P_i}{x_i}, \end{aligned} \quad (29)$$

and substitute Eq. (22), (25), and (29) into the LHS of Eq. (28), we obtain

$$\begin{aligned} &-\frac{1}{U} \cdot \frac{\partial U}{\partial \mathbf{x}} \cdot \left(\frac{\partial F}{\partial \mathbf{x}} \right)^{-1} \cdot \frac{\partial F}{\partial \beta} \\ &= \left[\frac{w_i a_i}{\sum_{k=1}^n w_k x_k a_k} - \frac{P_i}{x_i} \right]_{i \in [n]}^\top \cdot \left[D^{-1} - \frac{D^{-1} \mathbf{y} \mathbf{z}^\top D^{-1}}{1 + \mathbf{z}^\top D^{-1} \mathbf{y}} \right] \cdot \left[\frac{P_i}{x_i} \cdot (1 - 2P_i) \cdot \left(w_i - \sum_{k=1}^n P_k w_k \right) \right]_{i \in [n]} \\ &= \frac{1}{T} \sum_{i=1}^n \delta_i (1 - 2P_i) (w_i - T)^2 - \frac{(\sum_{i=1}^n \delta_i (1 - 2P_i) (w_i - T))^2}{T(1 + \sum_{i=1}^n \delta_i (1 - 2P_i))}, \end{aligned} \quad (30)$$

where $T = \frac{\sum_{k=1}^n w_k x_k a_k}{\sum_{k=1}^n x_k a_k} = \sum_{i=1}^n P_i w_i$. Therefore, it suffices to prove

$$\sum_{i=1}^n \delta_i (1 - 2P_i) (w_i - T)^2 - \frac{(\sum_{i=1}^n \delta_i (1 - 2P_i) (w_i - T))^2}{1 + \sum_{i=1}^n \delta_i (1 - 2P_i)} > 0. \quad (31)$$

Since $1 - 2P_i > 0$, from Cauchy–Schwarz inequality it holds that

$$\begin{aligned} & \sum_{i=1}^n \delta_i (1 - 2P_i) (w_i - T)^2 \cdot \left(1 + \sum_{i=1}^n \delta_i (1 - 2P_i) \right) \\ &= \sum_{i=1}^n \delta_i (1 - 2P_i) (w_i - T)^2 \cdot \sum_{i=1}^n \delta_i (1 - 2P_i) + \sum_{i=1}^n \delta_i (1 - 2P_i) (w_i - T)^2 \\ &\geq \left(\sum_{i=1}^n \delta_i (1 - 2P_i) (w_i - T) \right)^2 + \sum_{i=1}^n \delta_i (1 - 2P_i) (w_i - T)^2 \\ &> \left(\sum_{i=1}^n \delta_i (1 - 2P_i) (w_i - T) \right)^2, \end{aligned}$$

where the last inequality holds because $\{w_i\}$ are not identical so there exists at least one $j \in [n]$ such that $\delta_j (1 - 2P_j) (w_j - T)^2 > 0$. Therefore, Eq. (31) holds and we have $\frac{d \ln U}{d\beta} > 0$.

The monotonicity of $V(\beta)$: Next we show $\frac{dV}{d\beta} < 0$. Since $\frac{\partial V}{\partial \beta} = 0$, this is equivalent to show

$$\frac{\partial V}{\partial \mathbf{x}} \cdot \left(\frac{\partial F}{\partial \mathbf{x}} \right)^{-1} \cdot \frac{\partial F}{\partial \beta} > 0. \quad (32)$$

Note that $\frac{dV}{dx_i} = 1$, we have

$$\begin{aligned} & \frac{\partial V}{\partial \mathbf{x}} \cdot \left(\frac{\partial F}{\partial \mathbf{x}} \right)^{-1} \cdot \frac{\partial F}{\partial \beta} \\ &= - [1, 1, \dots, 1] \cdot \left[D^{-1} - \frac{D^{-1} \mathbf{y} \mathbf{z}^\top D^{-1}}{1 + \mathbf{z}^\top D^{-1} \mathbf{y}} \right] \cdot \left[\frac{P_i}{x_i} \cdot (1 - 2P_i) \cdot \left(w_i - \sum_{k=1}^n P_k w_k \right) \right]_{i \in [n]} \\ &= - \sum_{i=1}^n \frac{\delta_i x_i (1 - 2P_i) (w_i - T)}{P_i} + \frac{1}{1 + \sum_{i=1}^n \delta_i (1 - 2P_i)} \sum_{i=1}^n \frac{\delta_i x_i (1 - 2P_i)}{P_i} \sum_{i=1}^n \delta_i (1 - 2P_i) (w_i - T). \end{aligned} \quad (33)$$

To show the RHS of Eq. (33) is positive, it suffices to show

$$\left(1 + \sum_{i=1}^n \delta_i (1 - 2P_i) \right) \sum_{i=1}^n \frac{\delta_i x_i (1 - 2P_i) (w_i - T)}{P_i} < \sum_{i=1}^n \frac{\delta_i x_i (1 - 2P_i)}{P_i} \sum_{i=1}^n \delta_i (1 - 2P_i) (w_i - T). \quad (34)$$

Plugin $T = \sum_{i=1}^n P_i w_i$ into Eq. (34), it is equivalent to show that

$$\left(1 + \sum_{i=1}^n \delta_i (1 - 2P_i) \right) \sum_{i=1}^n \frac{\delta_i x_i (1 - 2P_i) w_i}{P_i} - \sum_{i=1}^n P_i w_i \sum_{i=1}^n \frac{\delta_i x_i (1 - 2P_i)}{P_i} < \sum_{i=1}^n \frac{\delta_i x_i (1 - 2P_i)}{P_i} \sum_{i=1}^n \delta_i (1 - 2P_i) w_i,$$

and note that $1 = \sum_{i=1}^n P_i$, it is equivalent to show

$$\sum_{i=1}^n [\delta_i (1 - 2P_i) + P_i] \sum_{i=1}^n \frac{\delta_i x_i (1 - 2P_i) w_i}{P_i} < \sum_{i=1}^n \frac{\delta_i x_i (1 - 2P_i)}{P_i} \sum_{i=1}^n [\delta_i (1 - 2P_i) + P_i] w_i. \quad (35)$$

Let $a_i = \frac{\delta_i x_i (1-2P_i)}{P_i}$, $b_i = \delta_i (1-2P_i) + P_i$, then Eq. (35) is equivalent to

$$\sum_{i=1}^n b_i \cdot \sum_{i=1}^n a_i w_i < \sum_{i=1}^n a_i \cdot \sum_{i=1}^n b_i w_i \iff \sum_{i>j} (w_i - w_j) \left(\frac{b_i}{a_i} - \frac{b_j}{a_j} \right) > 0.$$

Next, without loss of generality we show that for any $1 \leq i < j \leq n$, if $w_i > w_j$ then it also holds that $\frac{b_i}{a_i} > \frac{b_j}{a_j}$ for sufficiently large β . In fact, from $P_i = \frac{x_i e^{\beta w_i}}{\sum_{k=1}^n x_k e^{\beta w_k}}$ and Eq. (23) we obtain

$$\begin{aligned} \frac{b_i}{a_i} \cdot \frac{a_j}{b_j} &= \frac{\delta_j (\delta_i (1-2P_i) + P_i) (1-2P_j)}{\delta_i (\delta_j (1-2P_j) + P_j) (1-2P_i)} \cdot e^{\beta(w_i - w_j)} \\ &= \frac{1 + P_i \delta_i / (1-2P_i)}{1 + P_j \delta_j / (1-2P_j)} \cdot e^{\beta(w_i - w_j)} \\ &= \frac{1 + \frac{P_i^3}{(1-2P_i)(P_i^2 + x_i^2 c_i''(x_i))}}{1 + \frac{P_j^3}{(1-2P_j)(P_j^2 + x_j^2 c_j''(x_j))}} \cdot e^{\beta(w_i - w_j)}. \end{aligned} \quad (36)$$

On the one hand, because $1-2P_i, 1-2P_j \in [\frac{1}{3}, 1]$, $x_i, x_j \in [0, L]$, and $c_i''(x_i), c_j''(x_j) > 0$, the first terms of the LHS of Eq. (36) is a positive number lower bounded away from zero. On the other hand, $w_i > w_j$ ensures that the second term $e^{\beta(w_i - w_j)}$ can be arbitrarily large as long as β is sufficiently large. Therefore, there must exist a $\beta_0 > 0$ such that for any $\beta > \beta_0$, $\frac{b_i}{a_i} \cdot \frac{a_j}{b_j} > 1$ holds for any $i > j$. As a result, $(w_i - w_j) \left(\frac{b_i}{a_i} - \frac{b_j}{a_j} \right) > 0$ holds, which completes the proof. \square

B Additional Experiments

We use the following Multi-agent Mirror Descent (MMD) algorithm as the PNE solver of C^4 , whose convergence is guaranteed by [7]. Since each creator's strategy set $\mathcal{X}_i = [0, +\infty)$, we can simply choose a projection mapping $\text{Proj}_{\mathcal{X}_i}(\mathbf{x}) = (\max(x_i, 0))_{i=1}^n$. The gradients of utility functions can be implemented directly since they have closed forms. Through our experiment, the default $T = 10000, \eta = 0.1, \epsilon = 1e-2, x_i^{(0)} = 1.0$. Algorithm 2 is a simplified version of Algorithm 1 in [7] where we replace the gradient estimation to the exact gradient. According to Theorem 5.1 in [7], Algorithm 2 converges to the unique PNE of any C^4 game with probability 1.

Algorithm 2: Multi-agent Mirror Descent (MMD) with perfect gradient

Input: Maximum iteration number T , step size η , each player i 's utility function u_i , error tolerance ϵ , initial strategy $\mathbf{x}_i = \mathbf{x}_i^{(0)}$.

```

1 for  $t \in [T]$  do
2   if  $\|(\mathbf{g}_1, \dots, \mathbf{g}_n)\|_2 < \epsilon$ , then
3     Break
4   Compute gradient  $\mathbf{g}_i = \nabla_i u_i(\mathbf{x}_i, \mathbf{x}_{-i})$ , for  $i \in [n]$ ,
5   Update  $\mathbf{x}_i \leftarrow \text{Proj}_{\mathcal{X}_i}(\mathbf{x}_i + \eta \mathbf{g}_i), \forall i \in [n]$ .
```

Output: $(\mathbf{x}_1, \dots, \mathbf{x}_n)$.

Figure 3 illustrates the trade-off between U and V in the MovieLens environment, and Figure 4 plots the same information as shown in Figure 2 but with a different value of $\lambda = 0.1$. As we can see, the optimal $\text{PPM}(\beta)$ found by Algorithm 1 conveys a consistent message: it prioritizes the recommendation accuracy for users with more determined preferences while increasing the exploration strength for less selective users.

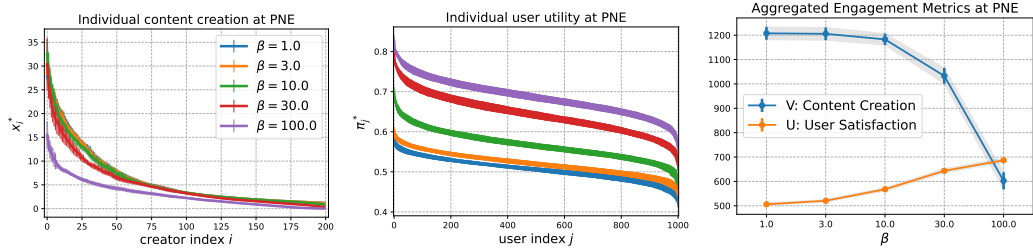


Figure 3: The left and the middle panel: the empirical distributions of content creation frequency x_i^* and each user's individual utility π_j^* . Different colors represent results for PNEs induced by different β . Right: the total content creation V and total user satisfaction U obtained under different β . Error bars obtained from 10 independently generated environments.

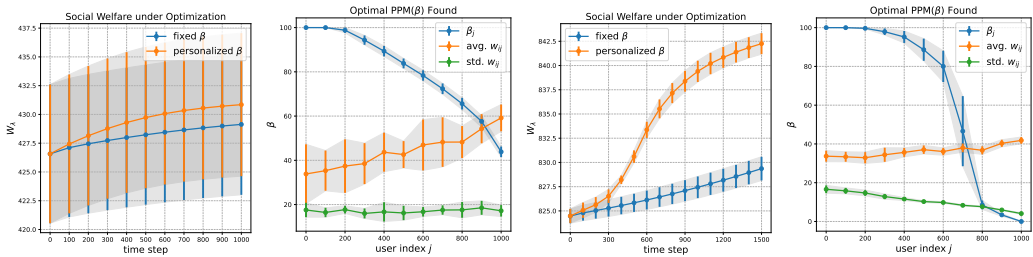


Figure 4: Panel 1,2: social welfare improving curve under Algorithm 1, and the distribution of the obtained optimal β_j in the synthetic environment. Panel 3,4: the same plots in the MovieLens environment. $\lambda = 0.1$.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: We list the paper's contribution and scope clearly at the last paragraph of Section 1

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We have discussed our limitations in the Conclusion section as well as in Theorem 2. We are not able to prove the most general claim with $m > 1$ and have to demonstrate our idea in a simplified setting ($m = 1$). However, our empirical study shows our theory is applicable to a much broader scenario.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: We have listed all the details of proof assumptions in the theory statement, and complete proofs in Appendix.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We provide the code and data in the supplementary material with a well commented python notebook for reproducing all the results and figures. Experiment details are listed in Section 7.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in

some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We provide the code and data in the supplementary material.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Experiment details are listed in Section 7.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We provide error bars in all figure in our experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Our experiments are not computational intense and do not require GPU as they are not related to deep learning.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: The codes are anonymous as well.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: It has been discussed at Abstract and Section 1.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.

- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.