

# On the Power of Strategic Corpus Enrichment in Content Creation Games

Haya Nachimovsky, Moshe Tennenholtz

Technion - Israel Institute of Technology, Haifa, Israel  
haya.nac@campus.technion.ac.il, moshet@technion.ac.il

## Abstract

Search and recommendation ecosystems exhibit competition among content creators. This competition has been tackled in a variety of game-theoretic frameworks. Content creators generate documents with the aim of being recommended by a content ranker for various information needs. In order for the ecosystem, modeled as a content ranking game, to be effective and maximize user welfare, it should guarantee stability, where stability is associated with the existence of pure Nash equilibrium in the corresponding game. Moreover, if the contents' ranking algorithm possesses a game in which any best-response learning dynamics of the content creators converge to equilibrium of high welfare, the system is considered highly attractive. However, as classical content ranking algorithms, employed by search and recommendation systems, rank documents by their distance to information needs, it has been shown that they fail to provide such stability properties. As a result, novel content ranking algorithms have been devised. In this work, we offer an alternative approach: corpus enrichment with a small set of fixed dummy documents. It turns out that, with the right design, such enrichment can lead to pure Nash equilibrium and even to the convergence of any best-response dynamics to a high welfare result, where we still employ the classical/current content ranking approach. We show two such corpus enrichment techniques with tight bounds on the number of documents needed to obtain the desired results. Interestingly, our study is a novel extension of Borel's Colonel Blotto game.

## 1 Introduction

Content creators aim that their content will be recommended by search and recommendation systems when faced with information needs expressed either implicitly or explicitly by the users (Kurland and Tennenholtz 2022). In abstract terms, a content creator (publisher) creates a content item (document) with the aim of being selected by a content ranking algorithm to be shown to a user with a particular topic of interest. At first approximation, a content creator aims that his document will be selected for as many topics (e.g., relevant queries) taken from a given set (Nachimovsky et al. 2024).

This general setup has been widely acknowledged by the AI community and various models have been advocated. In

the line of work we consider, content creators aim to promote their content rather than serve as malicious attackers in the system (as in, e.g., (Chen et al. 2023a,b; Liu et al. 2023)). Existing competitive content creation models emphasize the search ecosystem (Raifer et al. 2017) and the recommendation ecosystem (Ben-Porat and Tennenholtz 2018; Ben-Porat et al. 2018). Different models also vary in their representation of users and content, which can be explicit (Yao et al. 2023; Eilat and Rosenfeld 2023; Jagadeesan, Garg, and Steinhardt 2023; Hron et al. 2023) or via queries (Goren et al. 2021). Others discuss a more general notion of topics of interest (Ben-Basat, Tennenholtz, and Kurland 2017). However, all call for game-theoretic analysis of strategic content creation games, where publishers aim that some content ranking algorithm will select their content.

One property that is crucial for the above efforts is the general observation that in order to lead to high user social welfare, we need to care for publishers' utility (Mladenov et al. 2020), and that in order to maintain publishers, the game should converge to some stable deterministic equilibrium (Ben-Basat, Tennenholtz, and Kurland 2017; Yao et al. 2023; Ben-Porat and Tennenholtz 2018; Ben-Porat, Rosenberg, and Tennenholtz 2019). The related stability concept is captured by the requirement that the corresponding content ranking algorithm determines a game with pure Nash equilibrium, a deterministic content profile where unilateral deviations are not beneficial. Moreover, if that equilibrium can be shown to be obtained under best response learning dynamics, the system is especially appealing as any reasonable behavior (starting with any contents and with arbitrary order of modifications) will result in such a stable situation.

Given the above, a central theme of foundational work on strategic recommendation and search ecosystems originates from the fact that existing content ranking algorithms are unstable when facing such competition. Namely, given an information need, recommending the content closest to it (according to some well-crafted distance function estimating relevance) determines a game that does not possess a pure Nash equilibrium, hence might lead to harmful fluctuations. Indeed, most recent literature on strategic search and recommendation ecosystems attempts to offer alternative content ranking algorithms that do provide such stability (Yao et al. 2023; Ben-Porat and Tennenholtz 2018; Ben-Basat, Tennenholtz, and Kurland 2017).

One may wonder whether stability can be obtained if we do not modify the content ranking algorithms and stick to the well-crafted classical ones. In this work, we introduce and illustrate such general alternative: *strategic corpus enrichment*. To initiate the study of this approach, we use a model introduced in Nachimovsky et al. (2024), expanding on studies in Raifer et al. (2017)) and in Ben-Basat, Tennenholtz, and Kurland (2017).

Given a set of possible topics of interest (queries), each publisher needs to create a fixed-length document with content split among the topics. The content ranker provides a score for each publisher for each topic, where this score is a non-decreasing function of the weight the provider assigns to that topic (while this is typical for most existing systems, one can think about the proportion of news dedicated to the topic or the number of words associated with the corresponding query as naive estimates). For each topic, the publisher with the highest score is selected, with a uniform tie-breaking rule. A publisher strategy is the split of weights among topics, and its utility is the sum of topics it wins (Nachimovsky et al. 2024). While in recommendation systems the model is illustrative for current strategic content creation games, we observe that this model is an instance of one of the most celebrated economic models: the Colonel Blotto game, going back to Borel (1921). In this setting, there is a set of battlefields and each colonel needs to split a given fixed budget (army troops) between battlefields. A colonel wins a battlefield if his army is the largest, while aiming to win as many battlefields as possible. Hence, our content creation games are Colonel Blotto games among publishers, with queries as battlefields. Failures of the classical content ranking algorithm to obtain stable behavior are a direct implication of the lack of pure Nash equilibrium in Colonel Blotto games.

Nevertheless, in current AI ecosystems, we have a possibility that has not been considered so far: the (artificial) generation of static documents. If we add a relatively small number of static documents we modify the game. If this can be done in a way leading to pure Nash equilibrium where only the original documents win and (original) publishers even reach that equilibrium under any best-response learning dynamics, while obtaining high welfare, we can reach a stable and effective ecosystem while using the classical well-studied content ranking algorithm!

The above is exactly what we manage to obtain in our work. We begin by introducing the *uniform corpus extension technique* and show that it leads to our desired results. This approach reveals that irrespective of the number of topics or the classical content ranker employed, we can guarantee a pure Nash equilibrium and the convergence of any best-response learning dynamics to a welfare-maximizing result by at most doubling the number of documents. Furthermore, our findings establish tight bounds on the necessary number of documents under the problem parameters. Building on this foundation, we introduce the *generalized corpus extension technique*, which further reduces the number of documents required to achieve pure Nash equilibrium. However, in that context, best response dynamics do not always converge (so the system itself should recommend to publishers

how to split documents among topics to lead to stability). Our results underscore the effectiveness of using strategic corpus enrichment to achieve stability.

To summarize, our contributions are as follows:

- We propose the method of corpus enrichment as a novel strategy for mediators to influence publisher behavior within search and recommendation systems.
- We demonstrate that our approach consistently achieves stability (equilibrium), addressing the limitations of prior approaches where pure equilibria rarely exist. Furthermore, we provide tight bounds on the minimum number of documents needed in the corpus enrichment technique to ensure stability.
- Our equilibrium profiles not only ensure that publisher exposure remains unaffected but also maximize user social welfare, i.e., they maximize both publisher exposure and user social welfare.
- Our method guarantees convergence of best-response dynamics, while without corpus enrichment, they fail to converge despite the presence of a pure equilibrium. Additionally, we propose a computationally efficient algorithm for calculating the best response strategies.

## 2 Model

### 2.1 Multiple-Queries Ranking Game

Our model is based on the framework introduced by Nachimovsky et al. (2024) for ranking competitions between publishers aiming to be ranked highly for multiple queries. The *Multiple-Queries Ranking Game*  $G = \langle n, m, p \rangle$  is defined as a  $n$ -player’s game, where each player is a publisher who writes and modifies a single document. Let  $Q$  be a query set that contains  $m$  queries:  $Q = \{q_1, \dots, q_m\}$ . The term queries may refer to arbitrary content types or to different query aspects. A document  $d$  is defined as an element in  $D := \{(d^1, \dots, d^m) \in [0, 1]^m : \sum_{j=1}^m d^j \leq 1\}$ . Each component  $d^j$  ( $\in [0, 1]$ ) signifies the degree to which the document  $d$  focuses on the query  $q_j$ .

Let  $f : [0, 1] \rightarrow [0, 1]$  be a single peak function, meaning there exists a unique  $p \in [0, 1]$  such that  $f$  is non-decreasing in  $[0, p]$  and non-increasing in  $[p, 1]$ . The retrieval function  $r : D \times Q \rightarrow [0, 1]$  assigns each document  $d$  ( $\in D$ ) a score with respect to a query  $q_j$ :  $r(d, q_j) := f(d^j)$ . The documents are then ranked in descending order of scores (assuming ties are broken arbitrarily). Notice that if the function is constant in  $[p, 1]$  (or in particular,  $p = 1$ ) then we have here the classical selection of the document as most relevant to the query. We allow the function to be non-increasing in  $[p, 1]$  to deal with issues such as keyword stuffing (preventing spam by overloading query terms). None of our results will change given that assumption.

The document  $d_i$  published by player  $i$  is considered the pure *strategy* of player  $i$ . This strategy involves distributing the document’s focus across the various queries. A *strategy profile*  $s := (d_1, \dots, d_n)$  is the tuple of pure strategies of the players in the game. We use  $s_{-i} := (d_1, \dots, d_{i-1}, d_{i+1}, \dots, d_n)$  to denote the strategies of all players except player  $i$ .

The *utility* of a player  $i$  who writes document  $d_i$ , denoted  $U_i(s)$ , is defined as the number of queries where  $d_i$  achieves the highest rank among all documents. A strategy profile  $s$  is a (pure) *Nash equilibrium* if no player  $i$  has an incentive to change their strategy given the strategies  $s_{-i}$  of all other players. In other words, modifying their strategy does not lead to higher utility.

## 2.2 Corpus-Enriched Ranking Game

Adopting game-theoretic terminology, we will refer to the content selection mechanism (e.g. search engine, recommendation, system) as a mediator. The mediator’s goal is to promote the desired behavior of publishers, such as stability and satisfaction of the users. We propose a novel approach for the mediator to achieve this goal: adding a set of static documents  $T$  to the corpus. With these static documents, the rank of each of the publisher’s documents  $d_i$  for a query depends not only on the strategies of all other players but also on the static documents.

The *Corpus-Enriched Ranking Game* is defined as  $G_T = \langle n, m, p, T \rangle$ . In the game  $G_T$ , the rank of each of the publisher’s documents  $d_i$  for a query depends not only on the strategies of all other players but also on the static documents  $T$ .

Note that publishers derive utility solely from the highest-ranked document for a given query, as any lower ranking yields zero utility. Furthermore, publishers are indifferent to the identity of the highest-ranked document or any other documents that rank lower within the set  $T$ . We can therefore define a simplified version of the game  $G_T$ .

## 2.3 Simplified Corpus-Enriched Ranking Game

Let  $\vec{t}_j$  denote the value of the document in  $T$  corresponding to the highest ranking score with respect to query  $q_j \in Q$ . The impact of  $T$  on publishers’ behavior can be captured solely by  $\vec{t}$ , where  $\vec{t}_j$  essentially denotes a *threshold* that any publisher must surpass to achieve the highest rank for that query.

We define the *Ranking Game with Thresholds*  $G_{\vec{t}} = \langle n, m, p, \vec{t} \rangle$  with  $\vec{t} \in [0, 1]^m$ , where the static documents are represented using a single vector of thresholds. Thresholds in our model function similarly to reserved prices in auctions. In auctions, enforcing the reserved price does not necessitate physical bids but can be achieved by automatically rejecting any bids that fall below this threshold. Similarly, the mediator could set a publicly known threshold for each query, ensuring that a document cannot be considered the winner if its score is below this threshold.

**Lemma 2.1.** *For a given  $\eta \in \mathbb{N}$ , there exists a document set  $T$  such that  $|T| = \eta$  and the game  $G_T = \langle n, m, p, T \rangle$  has a pure equilibrium only if there exists a vector  $\vec{t} \in [0, 1]^m$  such that  $\|\vec{t}\|_1 \leq \eta$  and the game  $G_{\vec{t}} = \langle n, m, p, \vec{t} \rangle$  has a pure equilibrium.*

All proofs of lemmas, theorems, and corollaries presented in this paper are detailed in the supplementary material. The study of the simplified corpus enrichment game will determine the structure of the documents needed for effective corpus enrichment. This will be used to optimize the number of

documents needed to obtain equilibrium. Each such document corresponds to a subset of queries where the sum of thresholds of each query in the subset is no more than 1.

We specifically aim at corpus enrichment leading to an equilibrium in which only documents published by original publishers emerge as winners, that is, in all queries, there is at least one publisher who scores higher than all static documents. In the case of a tie between a publisher’s document and a static document, we assume that the mediator favors the document published by a real player, and ties between publishers’ documents are broken uniformly. Therefore, we require  $\max_{1 \leq i \leq n} r(d_i, q_j) \geq \vec{t}_j$  for every query  $q_j \in Q$ .

## 3 Preliminaries

The following theorem provides a full characterization of pure equilibrium existence in the game without corpus extension.

**Theorem 3.1** (Nachimovsky et al. (2024)).  *$G = \langle n, m, p \rangle$  has a pure Nash equilibrium iff  $p \leq \frac{1}{\max\{\lceil \frac{2-m}{n} \rceil, 1\}}$ .*

In the case where  $n < m$ , a pure equilibrium rarely exists. Specifically, a pure equilibrium exists only when the peak value is relatively small. Henceforth, we assume  $n < m$  since in the case where  $m \leq n$ , a pure equilibrium always exists.

The main idea of the proof of Theorem 3.1 for the case where  $n < m$  was to show that two properties are necessary for a pure equilibrium to exist: 1. the winner’s value must match the peak value across all queries 2. no player is the unique winner in more than one query.

**Notations** Given some strategy profile  $s$ , denote by  $w_j(s)$  the value of the document ranked highest in query  $q_j \in Q$ . The number of documents written by original players assigned the same highest retrieval score for query  $q_j$  when  $s$  is played is denoted by  $h_j(s)$ . A player  $i$  is called a winner in query  $q_j$  if  $t \leq d_i^j = w_j(s)$ . For every player  $i$ , denote by  $J_i(s)$  the set of queries where  $i$  is a winner, i.e.;  $w_i(s) = d_i^j$ . Let  $J_i^{solo}(s)$  and  $J_i^{tie}(s)$  denote the sets of queries where  $i$  wins uniquely and in a tie, respectively.

## 4 Uniform Corpus Enrichment

We begin our analysis by examining a subset of games with uniform  $\vec{t}_j$  across all queries, i.e., there exists some  $t \in [0, 1]$  such that  $\vec{t}_j = t$  for every  $q_j \in Q$ . With slight abuse of notation, we denote the game by  $G_t = \langle n, m, p, t \rangle$ ; and show that even within this constrained framework, it is possible to attain equilibrium by adding a relatively small number of static documents.

### 4.1 Equilibrium

Based on the results of Theorem 3.1, we identify the minimum threshold required to ensure the existence of an equilibrium where the winning value matches the peak for all queries. In this context, the threshold’s purpose is to prevent a player from being able to win in at least one more query.

**Theorem 4.1.**  $G_t = \langle n, m, p, t \rangle$  with  $\frac{1}{\lceil \frac{2m}{n} \rceil - 1} < p \leq \frac{1}{\lceil \frac{m}{n} \rceil}$  has a pure equilibrium iff

$$t > p - \frac{p \cdot \left( \left\lfloor \frac{1}{p} \right\rfloor + 1 \right) - 1}{\lceil \frac{2m}{n} \rceil - \left\lfloor \frac{1}{p} \right\rfloor}$$

Moreover, there exists a pure equilibrium  $s^*$  where  $w_j(s^*) = p$  for all  $q_j \in Q$ .

The proof idea is to show that the first property must still hold (regardless of  $t$ ), and then derive a bound on the threshold that will prevent a player from having a profitable deviation even when he is the unique winner in more than one query, i.e., the second property doesn't hold.

For larger peak values, achieving an arrangement in which the winners in all queries consistently reach the peak becomes impossible. As a result, any potential equilibrium will inevitably have at least one query in which the winning value is lower than the peak. This introduces additional complexities within the game.

Consider a scenario where a player is one of the winners in a particular query but not the sole winner. If the winning value for this query is below the peak value, even the smallest amount of excess resources enables the player to deviate profitably by becoming the sole winner of that query. In contrast, if the winning value of the query already equals the peak value, then no strategy would allow this player to become the sole winner in that query. Additionally, if a player is tied in three or more queries where the winning value is below the peak, they invariably have a profitable deviation available, where they can become the sole winner in all but one of these queries.

Despite these additional difficulties, our analysis reveals that there exists a specific threshold level at which achieving equilibrium remains feasible. Interestingly, this threshold remains constant and does not vary with  $p$  (except for the conditions that define its range).

**Theorem 4.2.**  $G_t = \langle n, m, p, t \rangle$  where  $\frac{1}{\lceil \frac{m}{n} \rceil} < p$  has a pure equilibrium iff there exists some  $\epsilon > 0$  such that

$$t \geq \begin{cases} \left(1 - \frac{n}{m}\right) \cdot \frac{n}{m} + \epsilon, & m \pmod n = 0 \\ \frac{1}{\lceil \frac{m}{n} \rceil}, & \text{otherwise} \end{cases}$$

Additionally, there is a pure equilibrium  $s^*$  where  $w_j(s^*) = \frac{1}{\lceil \frac{m}{n} \rceil}$  for all  $q_j \in Q$ .

In the case where the number of queries is exactly divisible by the number of players, there exists some arrangement where there are no ties between the players. Specifically, there are no ties where the winning value is lower than the peak. However, if  $m \pmod n > 0$ , it is inevitable to have at least one query with more than one winner. This situation necessitates additional constraints, as previously mentioned.

We found the minimal threshold required for the existence of an equilibrium for any number of players, queries, and peak value. It can be concluded that for any game  $G_t$ , the minimum number of documents required to reach equilibrium does not exceed the number of players. This finding holds true regardless of the number of queries, even as




Figure 1: Illustration of different equilibrium profiles in the game  $G_t = \langle 2, 5, 0.5, t \rangle$  for varying  $t$  values.

$m \rightarrow \infty$ , indicating that merely doubling the number of documents is sufficient to stabilize the system.

**Corollary 4.3.** Given any number of players  $n \in \mathbb{N}$ , number of queries  $m \in \mathbb{N}$  and peak value  $p \in (0, 1]$ , there exists a document set  $T$  such that  $|T| \leq n$ , for which the game  $G = \langle n, m, p, T \rangle$  has a pure equilibrium.

## 4.2 Social Welfare

The publishers' exposure in the equilibria we found is always maximized since only original publishers emerge as winners. Interestingly, while one might expect a tradeoff between the publishers' and users' welfare, we show that our proposed method of corpus enrichment can maximize both.

The selection of  $T$  influences not only the existence of an equilibrium but also its pattern. Figure 1 illustrates an example of a game where even a minimal increase in the threshold value substantially changes the equilibrium profile. In the first game (Figure 1a) the winning value for all queries is exactly  $\frac{1}{3}$ . In the second game (Figure 1b) the winners' scores increased in 4 out of 5 queries, which benefits user welfare, but no original publisher wins in  $q_3$ .

Typical semantics in the literature when discussing multiple queries is of several aspects of the same topic, and the aim is to provide good coverage for all different aspects. For example, in the context of the "Paris Olympic Games", we might analyze multiple queries addressing different aspects such as events, venues, and participants. Social welfare is defined as the distance from the peak value of the worst-performing query to emphasize coverage across all query aspects. Formally,

**Definition 1** (Social Welfare).

$$SW(s) = - \max_{q_j \in Q} |p - w_j(s)|$$

We show that the equilibria described in Theorems 4.1 and 4.2 not only minimize the number of documents that need to be added to achieve an equilibrium but also maximize the users' social welfare.

**Lemma 4.4.** The equilibria achieved in Theorems 4.1 and 4.2 maximize both the users' social welfare and publishers' exposure.




Figure 2: Possible fair and unfair improvement steps in the game  $G = \langle 2, 4, 1, 0.4 \rangle$ , with the player that deviates marked in red.

## 5 Best-Response Dynamics

Best-response dynamics refers to the process in which players update their strategies iteratively when they have a profitable deviation from their current strategy. Specifically, at each time step, an arbitrary player among those who can improve their outcome selects a strategy that maximizes their utility. These dynamics continue until a state of equilibrium is reached, where no player has a profitable deviation.

In the absence of external intervention, such as corpus enrichment, best response dynamics are known to potentially fail in converging (Nachimovsky et al. 2024), even when pure equilibrium exists. In this section, we prove that the best response dynamics do converge when using the corpus enrichment method for  $\frac{1}{\lfloor \frac{m}{n} \rfloor + 1} < p$ .

The *best response* is a strategy that yields the highest payoff for a player, given the strategies chosen by other players. Formally, a strategy  $d_i^l$  is deemed a *best response* of player  $i$  if it satisfies  $U_i(s_{-i}, d_i^l) \geq \max_{d_i \in D} U_i(s_{-i}, d_i)$ . A *best response improvement step* is a profile  $(s_{-i}, d_i^l)$ , where  $d_i^l$  is the best response of player  $i$ . Henceforth, we will refer simply to an improvement step when discussing best response improvement steps. An *improvement path* is defined as a sequence of such steps, denoted  $\{s^l\}_{l=0}$ . In each step  $l$ , exactly one player deviates from his action in the previous step,  $s^{l-1}$ , and this player is indicated by  $i_l$ . The players making these deviations may vary across different steps.

We introduce a mild assumption to address tie-breaking when players have multiple viable deviation strategies to select from.

**Definition 2** (Deviation Equity). For any player  $i$  and deviation  $d_i^l$ , the *deviation equity* is defined as the utility of

the worst-off player among all those whose utilities are decreased by the deviation. Formally,

$$v(d_i^l, s) = \min_{i': U_{i'}(d_i^l, s_{-i}) < U_{i'}(s)} U_{i'}(d_i^l, s_{-i})$$

If no player was harmed by  $i$ 's deviation to  $d_i^l$ ,  $v(d_i^l, s) = \infty$ .

**Definition 3** (Fair Player). A player  $i$  is *fair* if given multiple possible best response strategies, he selects the deviation  $d_i^{br}$  that maximizes the *deviation equity*.

The fairness assumption is made to avoid trivial divergence as the one depicted in Figure 2. In the example, two distinct improvement steps are considered, both resulting in the same utility gain for the deviating player. When the player opts for an "unfair" strategy, the resulting strategy profile (Figure 2b) merely permutes the initial strategy profile (Figure 2a), and therefore convergence is not guaranteed. In contrast, when a fair strategy is employed, the resulting strategy profile (Figure 2c) is a pure equilibrium.

Under the fairness assumption, if  $m \pmod n = 0$ , any best response deviation from a strategy by a player does not decrease the utility of others, leading to convergence in dynamics.

**Theorem 5.1.** Consider a game  $G = \langle n, m, p, t \rangle$  where  $m \pmod n = 0$  and  $p > \frac{1}{\lfloor \frac{m}{n} \rfloor + 1}$ . Any sequence of best response dynamics  $\{s^l\}_{l=0}$  where all players are fair players converges to a pure equilibrium  $s^*$  in no more than  $n$  steps.

To prove convergence in the general case, we proceed by identifying two potential patterns for the deviation of  $i_l$ :

- **type 1:**  $i_l$  deviates to a strategy in which he is the sole winner in  $\lfloor \frac{m}{n} \rfloor$  queries.
- **type 2:**  $i_l$  deviates to a strategy where he wins in  $\lfloor \frac{m}{n} \rfloor + 1$  queries, with at least  $\lfloor \frac{m}{n} \rfloor - 1$  of them uniquely.

In the following lemma we demonstrate that for any given strategy profile  $s$ , each player's best response strategy falls into one of these two patterns. Let  $x_i(s) = \{q_j \in Q : w_j(s_{-i}) < t\}$  be the set of queries in which, under  $s_{-i}$ , there is no document published by an original player whose score surpasses the threshold.

**Lemma 5.2.** Let  $G = \langle n, m, p, t \rangle$  with  $t \geq \frac{1}{\lfloor \frac{m}{n} \rfloor + 1}$ . Given any strategy profile  $s$ , the best response of each player  $i$  is either of type 1 or type 2. Specifically, if  $|x_i(s)| < \lfloor \frac{m}{n} \rfloor - 1$  then the best response is of type 1; if  $|x_i(s)| > \lfloor \frac{m}{n} \rfloor - 1$  it is of type 2; and if  $|x_i(s)| = \lfloor \frac{m}{n} \rfloor - 1$  it can be of either type.

Next, we demonstrate that if, during a sequence of fair best responses, a player deviates for the second time, then no other player that did not have a beneficial deviation at this point, will have a beneficial deviation in the next step.

**Lemma 5.3.** Let  $G = \langle n, m, p, t \rangle$  with  $t \geq \frac{1}{\lfloor \frac{m}{n} \rfloor}$ . If player  $i_l$  has previously deviated, then subsequent deviations by  $i_l$  ensure that no player who did not have a profitable deviation before will obtain one after this deviation.

The proof idea relies on the fact that after a first deviation, any subsequent deviations are of type 2, which satisfies this

property. Hence, a player's deviation can trigger future deviation by another player at most once, which is key to the main Theorem's proof.

**Theorem 5.4.** Let  $G = \langle n, m, p, t \rangle$  where  $m \pmod n > 0$ ,  $p > \frac{1}{\lfloor \frac{m}{n} \rfloor + 1}$  and  $t \geq \frac{1}{\lfloor \frac{m}{n} \rfloor}$ . If all players are fair, any sequence of best-response dynamics  $\{s^l\}_{l=0}$  will converge to a pure equilibrium  $s^*$  with rate  $O(\min\{m, n^2\})$ .

For lower peak values, best-response dynamics do not necessarily converge. Consider, for instance,  $G = \langle 2, 3, \frac{1}{2}, \frac{1}{2} - \epsilon \rangle$ , where  $\epsilon > 0$ . Starting from an initial strategy profile  $s_0 = ((\frac{1}{2} - \epsilon, 0, \frac{1}{2}), (0, \frac{1}{2}, \frac{1}{2} - \epsilon))$ , a possible deviation might lead to the profile  $s_1 = ((\frac{1}{2} - \epsilon, 0, \frac{1}{2}), (\frac{1}{2}, \frac{1}{2} - \epsilon, 0))$ . As  $s_1$  is a permutation of  $s_0$ , the dynamics result in a cycle.

### 5.1 Computationally Efficient Best Response Dynamics

The strategy space is continuous, so despite narrowing down the strategy patterns for any best-response deviation, there can still be infinitely many possible strategies. We therefore provide an efficient algorithm that calculates a best-response strategy, and show that using this algorithm, the fairness assumption can be relaxed, and that this relaxation is enough to obtain convergence.

---

Algorithm 1: (Almost) fair best response strategy construction for player  $i_i$

---

```

Initialize  $d_{br}^j = 0$  for all  $q_j \in Q$ .
Set  $d_{br}^j = t$  for all  $q_j \in x_{i_i}(s)$ .
if  $|x_{i_i}(s)| \geq \lfloor \frac{m}{n} \rfloor - 1$  then
    numQueriesToAdd =  $\lfloor \frac{m}{n} \rfloor + 1 - |x_{i_i}(s)|$ 
     $\epsilon = 0$ 
else
    numQueriesToAdd =  $\lfloor \frac{m}{n} \rfloor - |x_{i_i}(s)|$ 
     $\epsilon = \frac{1}{\lfloor \frac{m}{n} \rfloor (\lfloor \frac{m}{n} \rfloor + 1)}$ 
end if
for 1 to numQueriesToAdd do
     $i_{\max} = \arg \max_{i \neq i_i, |J_i^{i_i}(s_{-i_i}, d_{br})| \leq 1} |J_i(s_{-i_i}, d_{br})|$ 
     $q_{j_{\text{add}}} = \text{arbitrary element of } J_{i_{\max}}^{i_i}(s_{-i_i}, d_{br})$ 
     $d_{br}^{j_{\text{add}}} = t + \epsilon$ 
end for
return  $d_{br}$ 

```

---

Suppose a player can access the identity of original documents with highest score for each query, as well as their score. It is therefore easy to see that algorithm 1 runs in linear time and outputs a best-response strategy  $d_{i_i}^{alg}$ . It remains to show that this best-response strategy approximates deviation equity in a way that suffices to obtain our main result. This is guaranteed by the following Lemma:

**Lemma 5.5.** Let  $G = \langle n, m, p, t \rangle$  with  $t \geq \frac{1}{\lfloor \frac{m}{n} \rfloor}$ . If  $i_i$  performs a deviation of type 1, then

$$\left\lfloor \frac{m}{n} \right\rfloor - 1 < v(d_{i_i}^{alg}, s^l) \leq \max_{d_{i_i}^{br}} v(d_{i_i}^{br}, s^l)$$




Figure 3: Illustration of different equilibrium profiles in the game  $G_{\vec{t}} = \langle 2, 4, 1, \vec{t} \rangle$  for varying  $\|\vec{t}\|_1$  values.

and if the deviation is of type 2, then

$$\left\lfloor \frac{m}{n} \right\rfloor - \frac{1}{2} < v(d_{i_i}^{alg}, s^l) \leq \max_{d_{i_i}^{br}} v(d_{i_i}^{br}, s^l)$$

As it can be shown Lemma 5.5 implies Lemma 5.3 we get the desired result.

## 6 Generalized Corpus Enrichment

In the following section, we extend the results from Section 4 by allowing different thresholds for distinct queries. This relaxation allows for achieving equilibrium with fewer added documents. For example, let  $G_{\vec{t}} = \langle 4, 5, 1, \vec{t} \rangle$ . According to Theorem 4.2, an equilibrium profile in the game with uniform corpus extension will be

$$[s|\vec{t}] = \left[ \begin{array}{cccc|c} 0.5 & 0.5 & 0 & 0 & 0.5 \\ 0.5 & 0.5 & 0 & 0 & 0.5 \\ 0 & 0 & 0.5 & 0 & 0.5 \\ 0 & 0 & 0.5 & 0.5 & 0.5 \\ 0 & 0 & 0 & 0.5 & 0.5 \end{array} \right]$$

In the game above, using a uniform threshold requires the addition of three documents to reach equilibrium. However, considering that ties in query scores establish a "threshold" by themselves, we can optimize this by setting some scores to zero. The revised equilibrium profile, which minimizes the number of added documents, is as follows:

$$[s|\vec{t}] = \left[ \begin{array}{cccc|c} 0.5 & 0.5 & 0 & 0 & 0 \\ 0.5 & 0.5 & 0 & 0 & 0 \\ 0 & 0 & 0.5 & 0 & 0.5 \\ 0 & 0 & 0.5 & 0.5 & 0 \\ 0 & 0 & 0 & 0.5 & 0.5 \end{array} \right]$$

Employing non-uniform thresholds necessitates the addition of just one document, demonstrating a significant reduction in the number of dummy documents needed to achieve pure equilibrium existence. We proceed to analyze games with non-uniform thresholds.

### 6.1 Equilibrium

The following theorems are generalizations of the corresponding theorems from the uniform case, which we discussed in Section 4.

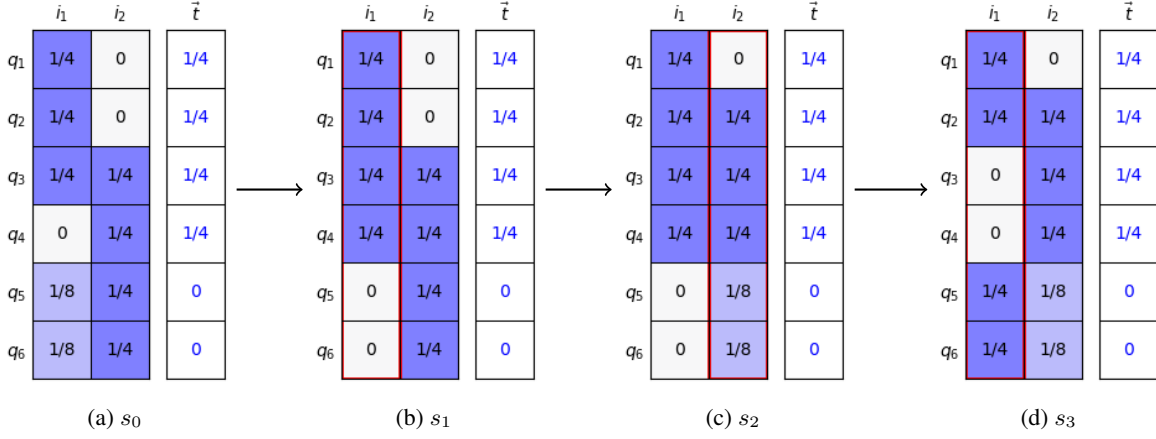


Figure 4: Improvement cycle in the game  $G = \langle 2, 6, 1, (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}, 0, 0) \rangle$ , with deviations marked in red at each step.

**Theorem 6.1.**  $G_{\vec{t}} = \langle n, m, p, \vec{t} \rangle$  with  $\frac{1}{\lfloor \frac{2m}{n} \rfloor - 1} < p \leq \frac{1}{\lfloor \frac{m}{n} \rfloor + 1}$  has a pure Nash equilibrium iff

$$\|\vec{t}\|_1 > n \cdot (1 - p) - 2p \cdot \left( \frac{1}{p} \right) - m$$

**Theorem 6.2.**  $G_{\vec{t}} = \langle n, m, p, \vec{t} \rangle$  where  $\frac{1}{\lfloor \frac{m}{n} \rfloor + 1} < p$  has a pure Nash equilibrium iff there exists some  $\epsilon > 0$  such that

$$\|\vec{t}\|_1 \geq \begin{cases} \frac{m-n}{\lfloor \frac{m}{n} \rfloor + 1}, & m \pmod{n} = 0 \\ \frac{m-n}{\lfloor \frac{m}{n} \rfloor} + \epsilon, & 1 \leq m \pmod{n} \leq n-2 \\ n \cdot \left( 1 - \frac{1}{\lfloor \frac{m}{n} \rfloor} \right) + \epsilon, & m \pmod{n} = n-1 \end{cases}$$

Since the generalized corpus extension method is a generalization of the uniform corpus extension method, we obtain:

**Corollary 6.3.** Let  $n \in \mathbb{N}$ . For any  $m \in \mathbb{N}$  and  $0 < p \leq 1$  there exists some  $\vec{t} \in \mathbb{R}^m$  such that  $\|\vec{t}\|_1 \leq n$  and the game  $G_{\vec{t}} = \langle n, m, p, \vec{t} \rangle$  has a pure equilibrium.

Comparing the threshold norms' bounds between uniform and generalized corpus extension methods shows that the difference in the lower bounds is most pronounced in games with a low query-to-player ratio. As this ratio increases, the difference diminishes, leading to the lower bounds on thresholds of the generalized method converging to those of the uniform method.

## 6.2 Social Welfare

In most instances, the equilibrium profile observed in the game with the general corpus extension method matches that obtained in the uniform setting. Nevertheless, there are specific situations where the equilibria differ, which impacts social welfare. Figure 3 illustrates an example of these differences by comparing two equilibrium profiles that arise under two different thresholds allocations, showing that social

welfare may not be maximized in equilibrium when thresholds are non-uniform. This comparison highlights the trade-off between minimizing the number of documents added to the corpus and maximizing social welfare.

## 6.3 Best-Response Dynamics

In Section 5 we showed that the game with uniform corpus extension has a desired property: best-response dynamics convergence. However, this property does not extend to the generalized method. Figure 4 presents an example from a game with general corpus enrichment featuring a non-uniform thresholds vector. As  $s_3$  and  $s_0$  are permutations of each other, the dynamics fail to converge.

## 7 Discussion

Our work introduced corpus enrichment as a technique for obtaining stable search and recommendation ecosystems. The resulting output maximizes social welfare, while only original content providers are selected, and the number of dummy documents needed is small. Moreover, we showed that for uniform corpus enrichment, this can be obtained by best-response dynamics. Generalized corpus enrichment allows to further restrict the number of dummy documents to be added to obtain stable output.

A limitation of our results is that there are restrictions on the settings in which best-response dynamics converge. In particular, best-response dynamics does not always converge under the generalized corpus enrichment technique. While as discussed in Section 5 our results can be obtained under some minimal observability assumptions of winning documents and their score, it is still a significant limitation, which may be of interest to further relax. Further work may consider some restrictions on the best-response dynamics beyond the notion of fairness that may suffice for convergence. One other possible extension of our study is to consider corpus enrichment strategies in which social welfare does increase due to the introduction of few high-quality documents while relaxing the requirement that only content written by the original authors is selected.

## Acknowledgements

This work was supported by funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement 740435).

## References

- Ben-Basat, R.; Tennenholtz, M.; and Kurland, O. 2017. A game theoretic analysis of the adversarial retrieval setting. *Journal of Artificial Intelligence Research*, 60: 1127–1164.
- Ben-Porat, O.; Goren, G.; Rosenberg, I.; and Tennenholtz, M. 2018. From Recommendation Systems to Facility Location Games. In *AAAI Conference on Artificial Intelligence*. AAAI Press. ISBN 978-1-57735-809-1.
- Ben-Porat, O.; Rosenberg, I.; and Tennenholtz, M. 2019. Convergence of Learning Dynamics in Information Retrieval Games. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01): 1780–1787.
- Ben-Porat, O.; and Tennenholtz, M. 2018. A game-theoretic approach to recommendation systems with strategic content providers. *Advances in Neural Information Processing Systems*, 31.
- Borel, E. 1921. La théorie du jeu et les équations intégrales à noyau symétrique. *Comptes rendus de l’Académie des Sciences*, 173 (1304-1308):58.
- Chen, X.; He, B.; Sun, L.; and Sun, Y. 2023a. Defense of Adversarial Ranking Attack in Text Retrieval: Benchmark and Baseline via Detection. *CoRR*, abs/2307.16816.
- Chen, X.; He, B.; Ye, Z.; Sun, L.; and Sun, Y. 2023b. Towards Imperceptible Document Manipulations against Neural Ranking Models. In *Findings of the Association for Computational Linguistics: ACL*, 6648–6664.
- Eilat, I.; and Rosenfeld, N. 2023. Performative Recommendation: Diversifying Content via Strategic Incentives. In Krause, A.; Brunskill, E.; Cho, K.; Engelhardt, B.; Sabato, S.; and Scarlett, J., eds., *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, 9082–9103. PMLR.
- Goren, G.; Kurland, O.; Tennenholtz, M.; and Raiber, F. 2021. Driving the Herd: Search Engines as Content Influencers. In *Proceedings of the 30th ACM International Conference on Information and Knowledge Management*. ACM.
- Hron, J.; Krauth, K.; Jordan, M. I.; Kilbertus, N.; and Dean, S. 2023. Modeling content creator incentives on algorithm-curated platforms. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Jagadeesan, M.; Garg, N.; and Steinhardt, J. 2023. Supply-Side Equilibria in Recommender Systems. In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- Kurland, O.; and Tennenholtz, M. 2022. Competitive search. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2838–2849.
- Liu, Y.; Zhang, R.; Guo, J.; de Rijke, M.; Chen, W.; Fan, Y.; and Cheng, X. 2023. Topic-oriented Adversarial Attacks against Black-box Neural Ranking Models. In *Proceedings of SIGIR*, 1700–1709.
- Mladenov, M.; Creager, E.; Ben-Porat, O.; Swersky, K.; Zemel, R.; and Boutilier, C. 2020. Optimizing long-term social welfare in recommender systems: A constrained matching approach. In *International Conference on Machine Learning*, 6987–6998. PMLR.
- Nachimovsky, H.; Tennenholtz, M.; Raiber, F.; and Kurland, O. 2024. Ranking-Incentivized Document Manipulations for Multiple Queries. In *Proceedings of the 2024 ACM SIGIR International Conference on Theory of Information Retrieval, ICTIR ’24*, 61–70. New York, NY, USA: Association for Computing Machinery. ISBN 9798400706813.
- Raifer, N.; Raiber, F.; Tennenholtz, M.; and Kurland, O. 2017. Information retrieval meets game theory: The ranking competition between documents’ authors. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 465–474.
- Yao, F.; Li, C.; Sankaraman, K. A.; Liao, Y.; Zhu, Y.; Wang, Q.; Wang, H.; and Xu, H. 2023. Rethinking Incentives in Recommender Systems: Are Monotone Rewards Always Beneficial? In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.