

The Search for Stability: Learning Dynamics of Strategic Publishers with Initial Documents

OMER MADMON^{*}, Technion - Israel Institute of Technology, Israel

IDAN PIPANO[†], Technion - Israel Institute of Technology, Israel

ITAMAR REINMAN[‡], Technion - Israel Institute of Technology, Israel

MOSHE TENNENHOLTZ[§], Technion - Israel Institute of Technology, Israel

We study a game-theoretic information retrieval model in which strategic publishers aim to maximize their chances of being ranked first by the search engine while maintaining the integrity of their original documents. We show that the commonly used Probability Ranking Principle (PRP) ranking scheme results in an unstable environment where games often fail to reach pure Nash equilibrium. We propose two families of ranking functions that do not adhere to the PRP. We provide both theoretical and empirical evidence that these methods lead to a stable search ecosystem, by providing positive results on the learning dynamics convergence. We also define the publishers' and users' welfare, demonstrate a possible publisher-user trade-off, and provide means for a search system designer to control it. Finally, we show how instability harms long-term users' welfare.

JAIR Associate Editor: Haris Aziz

JAIR Reference Format:

Omer Madmon, Idan Pipano, Itamar Reinman, and Moshe Tennenholtz. 2025. The Search for Stability: Learning Dynamics of Strategic Publishers with Initial Documents. *Journal of Artificial Intelligence Research* 83, Article 15 (July 2025), 36 pages. DOI: [10.1613/jair.1.17997](https://doi.org/10.1613/jair.1.17997)

1 Introduction

The rapid advancement of information retrieval systems, exemplified by search engines such as Google and Bing, and recommendation systems such as YouTube and Spotify, has reshaped how users interact with the vast digital information landscape. Central to their function is ad-hoc retrieval: ranking documents within a corpus according to their relevance to a user's query. A key phenomenon in such ecosystems is strategic behaviors from content providers, who modify their web content to improve their ranking in search results, a practice known as search engine optimization (SEO). Such strategies significantly influence the quality of search results and the user's experience. This phenomenon has garnered substantial attention, and game-theoretic tools have been applied to analyze it. Since a ranking scheme induces a game, the task of choosing a ranking scheme is essentially a task of mechanism design, a central field in game theory.

^{*}Equal Contribution.

[†]Equal Contribution.

[‡]Equal Contribution. Corresponding Author.

[§]Equal Contribution.

Authors' Contact Information: Omer Madmon, omermadmon@campus.technion.ac.il, Technion - Israel Institute of Technology, Haifa, Israel; Idan Pipano, idan.pipano@campus.technion.ac.il, Technion - Israel Institute of Technology, Haifa, Israel; Itamar Reinman, itamarr@campus.technion.ac.il, Technion - Israel Institute of Technology, Haifa, Israel; Moshe Tennenholtz, moshet@technion.ac.il, Technion - Israel Institute of Technology, Haifa, Israel.



This work is licensed under a [Creative Commons Attribution International 4.0 License](https://creativecommons.org/licenses/by/4.0/).

© 2025 Copyright held by the owner/author(s).

DOI: [10.1613/jair.1.17997](https://doi.org/10.1613/jair.1.17997)

One key challenge regarding these ranking-induced games is the design of *stable ranking mechanisms*, under which natural dynamics of content providers (hereafter, publishers) converge to a stable state, in which no publisher is willing to modify their content given the other publishers' content. This stable state is widely known as *Nash equilibrium*, and one popular family of learning dynamics is the class of *better response dynamics*, which assumes that at each round one publisher modifies her content in a way that increases her utility compared to the previous round. Note that this minimal assumption captures a wide range of behavioral patterns that are likely to be observed in practice. A search engine designer is typically interested in a ranking scheme that guarantees the convergence of actual dynamics induced by strategic content modification in the text space. While the structure of such dynamics is unknown, a stronger approach we adopt is to design a ranking scheme that guarantees the convergence of *any* better response dynamics. This approach in particular ensures the convergence of any real-life dynamics.

In this work, we study stable mechanisms for search and exposure-based recommendation systems in a continuous action space model. We adopt a framework that models the strategy space of publishers (the documents they create) as a continuous embedding space, distinguishes between publishers' and users' welfare, and accounts for publishers' initial documents, which can be thought of as preferred content that the publishers aim to maintain some level of integrity to. An important component of the model is the *information need*, which is the user's underlying goal or requirement for information. We assume a dense retrieval model, in which the information need is represented in the same continuous embedding space as the publishers' documents [29].

Within this model, we study the convergence of better response dynamics under different ranking schemes. Importantly, while we use the terminology of a search engine that determines publishers' exposure through an underlying ranking mechanism, our approach is agnostic to the specific ranking process. Instead, we assume that the system can directly allocate exposure to each publisher. This flexibility allows our model to extend beyond the context of information retrieval, making it applicable to exposure-based recommendation systems more broadly.

The **probability ranking principle (PRP)** is a widely used principle that states that documents should be ranked in order of decreasing probability of relevance to the user [38]. In our model, this translates into the principle of greedily allocating the entire exposure mass to the publisher who provides the most relevant content (in terms of distance from the information need). We begin our analysis by providing a negative result on the stability of the PRP ranking scheme (Observation 1). We then introduce two alternative ranking schemes, the **softmax ranking function** (also discussed in Hron et al. [21]) and the **linear ranking function**.

The softmax ranking function serves as a smoothed version of the PRP ranking function, with smoothing controlled by an inverse temperature hyperparameter. We observe that the softmax ranking function is part of a broad and natural class of *proportional ranking functions*. However, we also establish that none of these ranking functions induce games with an exact potential, a primary tool for proving and analyzing the convergence of better response dynamics (Theorem 2). While the absence of potential is not necessarily enough to preclude convergence of better-response dynamics, we also show an instance of the publishers' game in which softmax ranking can induce cyclic content modification patterns, thereby preventing convergence (Observation 2).¹

This contrasts with the linear ranking function, which applies a linear transformation (with the slope as a hyperparameter) to the publishers' *linear relative relevance*—defined as the difference between a publisher's proximity to the information need and the average proximity of their competitors. For this family of ranking functions, we prove the convergence of *any* better response dynamics to a stable state using a potential function argument (Theorem 3). We also provide a detailed equilibrium characterization and welfare-in-equilibrium analysis under the linear ranking function, utilizing its convenient theoretical properties (Theorem 4, Lemma 3

¹Throughout the paper, we use the term “potential” to refer to an *exact potential*, which is indeed only a sufficient condition for proving better response dynamics convergence. In general games, an exact potential might not exist, and better response dynamics will still be guaranteed to converge (for instance, there might exist a *generalized ordinal potential*, a weaker potential notion that also guarantees convergence of learning dynamics). We show this is not the case when using the softmax ranking function.

and additional discussion in Appendices A and B). Lastly, we show that it is impossible to design a valid ranking function by taking a non-linear transformation of the linear relative relevance, which justifies our design choice of linearity (Theorem 5).

Our theoretical analysis is followed by a rigorous empirical evaluation of all three ranking functions in terms of users' and publishers' welfare.² We perform comparisons both between ranking families and within ranking families. Our empirical analysis demonstrates how the instability of the PRP ranking function harms its performance compared to the alternatives we suggest. While indeed there are no theoretical guarantees for the softmax function, we empirically show it is stable—in the sense that learning dynamics converge in almost any instance—and hence manages to surpass the PRP in terms of users' welfare. This underscores the practical effectiveness of the softmax ranking function. Additionally, we show an inherent trade-off between publishers' and users' welfare and provide a means for a system designer to control it through hyperparameter tuning.

Paper organization. In §2 we review related work in the field of strategic information retrieval. §3 provides preliminary definitions and results from game theory. In §4 we discuss the publishers' game model. §5 provides a theoretical analysis of learning dynamics in our model, and studies stability under different ranking schemes (PRP, softmax, and linear). In §6 we provide experimental results of simulating learning dynamics under the different ranking schemes. We then conclude and present future work directions in §7. The Appendix includes further theoretical developments, additional empirical results, and proof segments omitted from the main article.

2 Related Work

Our work lies in the intersection of game theory, economics, and machine learning. There is extensive literature on strategic and societal aspects of machine learning, such as classification and segmentation of strategic individuals [19, 35, 25, 33], fairness in machine learning [14, 1, 9], and recommendation systems design [10, 7, 4, 30]. Specifically, we contribute to the study of learning dynamics within machine learning, data-driven environments, and ad auctions [16, 18, 28, 13].

In this work, we study a game-theoretic model of the search (or more generally, recommendation) ecosystem, in which strategic content providers compete to maximize impressions and visibility. Ben Basat et al. [6] and Ben Basat et al. [5] showed that the widely used PRP ranking scheme of Robertson [38] leads to sub-optimal social welfare, and demonstrated how introducing randomization into the ranking function can improve social welfare in equilibrium. Raifer et al. [37] provided both theoretical and empirical analysis of a repeated game, and revealed a key strategy of publishers: gradually changing their documents to resemble those ranked highest in previous rounds. Ben-Porat et al. [8] studies better-response dynamics in a publishers' game in which the strategy space is modeled as a finite set of "topics". Kurland and Tennenholtz [24] advocates for the devising of ranking functions that are optimized not only for short-term relevance effectiveness but also for long-term corpus effects.

While these works rely on representing documents as a finite set of "topics", a more recent line of work takes a different modeling assumption of representing content as a continuous, high-dimensional space. Hron et al. [21] studies both PRP and softmax-based ranking functions, but does not incorporate initial documents, which are crucial in our analysis. Jagadeesan et al. [22] studies a model similar to Hron et al. [21] but also incorporates quality constraints—the publisher's cost is a function of the document's norm. This is a private case of our model, in which all publishers have the same initial document. Our model generalizes by assuming publishers' optimal documents may be different, which is a rather realistic scenario.

In a similar spirit, Yao et al. [41] studies competition between human content creators and Generative AI, and Yao et al. [43] provides novel techniques for optimizing user welfare in equilibrium. Importantly, none of these papers study learning dynamics or their convergence. Closer to our work, Yao et al. [42] studies the design of stable recommendation mechanisms under similar modeling analysis as in our paper. However, their work

²Code and results are available in <https://github.com/ireinman/The-Search-for-Stability>.

concerns the design of a reward mechanism (which can be seen as designing the ranking and payments jointly), while in our model, the designer can only determine the exposure of each publisher, without being able to transfer money. Yao et al. [40] and Madmon et al. [27] also studied content provider dynamics, and adopted modeling assumptions similar to ours. However, their work focuses on regret minimization dynamics, where publishers act simultaneously and aim to minimize their cumulative regret, defined as the gap between the utility they could have achieved with full knowledge of others' actions and the utility they actually obtain. In contrast, we study better-response dynamics, where at each time step, a single publisher updates its content to myopically improve its own utility.

3 Preliminaries

In this section, we provide some standard definitions and results from the field of game theory. We start with the definition of a **game**:

DEFINITION 1. A game is a tuple $(N, (X_i)_{i \in N}, (u_i)_{i \in N})$, where $N := \{1, \dots, n\}$ is the set of players, and for each player $i \in N$, X_i is the set of actions. The set $X := X_1 \times \dots \times X_n$ is called the set of pure strategy profiles. Each player $i \in N$ has a utility function $u_i : X \rightarrow \mathbb{R}$.

In the sequel, we assume that X_i is a compact and convex set and that u_i is a bounded function. We also assume that all players are rational, in the sense that each player attempts to maximize her own utility. Each player's utility depends not only on her action but also on the actions taken by all other players. We denote by $X_{-i} := X_1 \times \dots \times X_{i-1} \times X_{i+1} \times \dots \times X_n$ the set of all pure strategy profiles of all players but player i , and use the notation of $x_{-i} \in X_{-i}$ to denote such a strategy profile of everyone but i . An important concept is the concept of **ε -better response**.

DEFINITION 2. Given a strategy profile $x_{-i} \in X_{-i}$ of all players but i and actions $x_i, x'_i \in X_i$ of player i , we call x'_i an ε -better response than x_i with respect to x_{-i} if $u_i(x'_i, x_{-i}) > u_i(x_i, x_{-i}) + \varepsilon$.

The notion of ε -best response allows defining the solution concept of an **ε -pure Nash equilibrium**: a pure strategy profile from which no player can improve her utility by more than ε by deviating (i.e., changing her actions given all other players' actions are fixed).

DEFINITION 3. A pure strategy profile $x \in X$ is called an ε -pure Nash equilibrium (ε -PNE), if for every player $i \in N$, **no** x'_i is an ε -better response than x_i with respect to x_{-i} .

For $\varepsilon = 0$, 0-better response and 0-PNE are called better response and PNE, respectively. A PNE is often considered a stable state of a game [10]. When players apply the concept of ε -better response we get the following type of dynamics:

DEFINITION 4. An **ε -better response dynamics** is a sequence of profiles $\{x^{(t)}\}_{t=1}^{\infty} \subseteq X$ such that in each timestep t , either $x^{(t-1)}$ is an ε -PNE and $x^{(t)} = x^{(t-1)}$, or there is a single deviator $i \in N$ such that $x_{-i}^{(t)} = x_{-i}^{(t-1)}$ and $x_i^{(t)}$ is an ε -better response than $x_i^{(t-1)}$ with respect to $x_{-i}^{(t-1)}$. We say such a dynamics has **converged** if there exists a timestep T such that $x^{(T)}$ is an ε -PNE.

Another class of games, which will be of interest in this paper, is the class of **potential games**.

DEFINITION 5 (MONDERER AND SHAPLEY). A game G is called a potential game if there exists a potential $\phi : X \rightarrow \mathbb{R}$, such that for any $i \in N, x_{-i} \in X_{-i}, x'_i, x''_i \in X_i$:

$$\phi(x'_i, x_{-i}) - \phi(x''_i, x_{-i}) = u_i(x'_i, x_{-i}) - u_i(x''_i, x_{-i}) \quad (1)$$

That is, in a potential game, a single function ϕ “captures” the incentives of all the players. It is straightforward that any potential maximizer is a PNE. Moreover, Monderer and Shapley [31] shows a significant result:

THEOREM 1 (MONDERER AND SHAPLEY). *Let G be a potential game and let $\varepsilon > 0$. Then, G possesses at least one ε -PNE, and any ε -better response dynamics in G converges.*

Another important concept is the concept of a **strictly dominant strategy**, a strategy which is the best choice for a player regardless of the strategies chosen by other players. The significance of strictly dominant strategies lies in their stability and predictability.

DEFINITION 6. *A strategy $x_i \in X_i$ is a strictly dominant strategy for player $i \in N$ if $\forall x_{-i} \in X_{-i}, x'_i \in X_i$, $u_i(x_i, x_{-i}) > u_i(x'_i, x_{-i})$.*

4 The Model

We start by defining the model of a publishers’ game. A publishers’ game consists of n publishers (players), each can provide content in the continuous k dimensional embedding space $[0, 1]^k$, where the distance between documents is measured by some continuous semi-metric d .³ This could be, for example, the Euclidean distance, which is widely used in the context of embedding [39, 2, 3].

Each publisher i has a preferred content type (or: initial document), denoted by $x_0^i \in [0, 1]^k$. We denote by $x^* \in [0, 1]^k$ the information need representation in the feature space. The term x^* can refer to a particular query or topic in the search environment.⁴ In our model, the information need is represented in the same embedding space as the documents, inspired by BERT rankers and dense retrieval models [36, 45, 46]. While the exact information need is typically latent, it is standard in the information retrieval literature to assume that it can be modeled as a point in a semantic or topical space [15].

We assume the information need x^* is known to the publishers. While this assumption may seem restrictive, we claim that it is justified, for example, in the setting of a *niche domain*, where typically a small number of publishers compete over a specific domain. In such scenarios, the publishers specialize in publishing content within a specific domain and know precisely the users’ preferences and needs. Examples for such domains include *travel content*, where publishers generate destination-specific articles tailored to well-understood traveler interests, and *professional content creators* who produce material for clearly defined audiences in domains in which the publishers are often experts themselves (such as finance or sports).

The search engine uses a ranking function r to determine the documents’ ranking. We assume that r is oblivious, meaning it is independent of the publishers’ identities. The choice of r induces a strategic interaction among the publishers, which can be formally defined as a game:

DEFINITION 7. *A **publishers’ game** G is a tuple $(N, k, \lambda, d, (x_0^i)_{i \in N}, x^*, r)$, where:*

- $N := \{1, 2, \dots, n\}$ is the set of publishers, with $n \geq 2$;
- k is the embedding-space dimension of documents;
- $X_i := [0, 1]^k$ is the set of publisher i ’s possible actions;
- $\lambda > 0$ is the factor of the cost for providing content that is different from the initial document;
- $d : [0, 1]^k \times [0, 1]^k \rightarrow \mathbb{R}_+$ ⁵ is a continuous semi-metric satisfying $\forall x, y : d(x, y) \leq 1$,⁶

³A semi-metric is a function that satisfies all the axioms of a metric, with the possible exception of the triangle inequality.

⁴Alternatively, when considering a recommender system that places items (documents) and users in the same latent space, x^* can be thought of as a representative user.

⁵We use \mathbb{R}_+ to denote the set of non-negative real numbers and \mathbb{R}_{++} to denote the set of strictly positive real numbers.

⁶This is without loss of generality, as for any distance function that does not satisfy $\max_{x, y} d(x, y) \leq 1$, one can simply divide it by the maximal distance.

- $x_0^i \in X_i$ is publisher i 's initial document;
- $x^* \in [0, 1]^k$ is the information need;
- $r : X \rightarrow \Delta^n$ is a ranking function, that might depend on the information need x^* .

The ranking function r determines the distribution over the winning publishers given a strategy profile $x \in X$, with the interpretation that $r_i(x)$ is the probability of publisher i being ranked first by r under the strategy profile x . This modeling assumption is motivated by the empirical evidence that users usually prefer to reformulate their query rather than searching beyond the first retrieved result [12, 26, 23]. We abuse the notation and denote $d^*(x_i) := d(x_i, x^*)$ and $d_i^0(x_i) := d(x_i, x_0^i)$. The utility of publisher i under the profile x is defined to be:

$$u_i(x) := r_i(x) - \lambda d_i^0(x_i) \quad (2)$$

That is, the utility of publisher i is the probability of winning minus the cost of providing content that differs from her initial document. We now define the **publishers' welfare** $\mathcal{U}(x)$ and the **users' welfare** $\mathcal{V}(x)$, which will henceforth be referred to collectively as **welfare measures**. The publishers' welfare is the sum of the publishers' utilities, while the users' welfare is the expected relevance of the first-ranked document, where we quantify the relevance of a document x_i by $1 - d^*(x_i)$.

$$\mathcal{U}(x) := \sum_{i=1}^n u_i(x) = 1 - \lambda \sum_{i=1}^n d_i^0(x_i) \quad (3)$$

$$\mathcal{V}(x) := \sum_{i=1}^n (1 - d^*(x_i)) \cdot r_i(x) = 1 - \sum_{i=1}^n d^*(x_i) \cdot r_i(x) \quad (4)$$

EXAMPLE 1 (RUNNING EXAMPLE). *To illustrate our model, we consider a simple example involving two publishers competing for exposure in response to a specific information need. In this toy example, the embedding space is two-dimensional ($k = 2$), and the distance metric is the squared Euclidean distance, defined as $d(x, y) = \frac{1}{2} \|x - y\|_2^2$.*

The publishers compete over a population of users whose information need can be described as adventurous eco-tourism. Let $x^ = (0.358, 0.908)$ be the embedding representation of the information need. While many small independent publishers may produce content relevant to this topic, we assume that, within this niche domain, there are only two prominent publishers, receiving the vast majority of exposure, with all other publishers receiving negligible attention in comparison. We therefore only consider the two prominent publishers in this running example.*

We assume that both publishers share a cost parameter of $\lambda = 1$. Publisher 1 specializes in Arctic wildlife and her initial document embedding vector is given by $x_1^0 = (0.950, 0.035)$, while Publisher 2 focuses on tropical hiking trails, having $x_2^0 = (0.933, 0.773)$ as her initial document representation.

We will refer to this running example throughout Section 5 to illustrate the behaviors induced by different ranking functions, highlighting properties such as the existence and characterization of PNE, as well as the convergence (or lack thereof) of learning dynamics.

5 Ranking Functions and Learning Dynamics

This section presents a theoretical analysis of the learning dynamics in our model and studies stability under three ranking functions: the PRP ranking function, the softmax ranking function, and the linear ranking function.

5.1 The PRP Ranking Function

The Probability Ranking Principle (PRP; [38]) is a key principle in information retrieval, which forms the theoretical foundations for probabilistic information retrieval. The principle says that documents should be ranked in decreasing order of the probability of relevance to the information need expressed by the query. In

our model, the analogous principle means that the closest document to the information need should always be ranked first by the ranking function. We begin by formally defining the PRP ranking function in our model:

DEFINITION 8. *The **PRP ranking function**, denoted by r^* , is the ranking function defined as follows:*

$$r_i^*(x) := \frac{1}{|\mu(x)|} \cdot \mathbf{1}_{i \in \mu(x)} \quad (5)$$

where $\mu(x) := \operatorname{argmin}_{i \in N} d^*(x_i)$ is the set of all publishers providing the closest content to the information need x^* .

Note that the PRP yields optimal short-term retrieval in the sense that it maximizes the users' welfare $\mathcal{V}(x)$ for a fixed strategy profile x . However, its discontinuity induces unstable ecosystems, in which there might not exist a PNE. This instability causes sub-optimal long-term effects, as we show later.

EXAMPLE 2. *Consider the instance presented in Example 1, with the PRP as its ranking function. Now, consider the strategy profile in which $x_1 = x_2 = x^*$, namely, both publishers precisely target the most suitable content given the information need. Under this profile, the utility of Publisher 1 is given by:*

$$\begin{aligned} u_1(x_1, x_2) &= r_1(x_1, x_2) - \lambda d_1^0(x_1) \\ &= r_1(x^*, x^*) - 1 \cdot \frac{1}{2} \|(0.950, 0.035) - (0.358, 0.908)\|^2 \\ &\approx \frac{1}{2} - 0.556 = -0.056 < 0 \end{aligned} \quad (6)$$

Now, observe that if Publisher 1 deviates and chooses $x_1 = x_0^1$, her utility becomes 0: on the one hand, she will never be ranked first, but on the other hand, she incurs no cost for deviating from her initial document. Therefore, the profile of $x_1 = x_2 = x^*$ is not a PNE, as Publisher 1 has a beneficial deviation.

In Example 2, we showed that the profile in which both players choose x^* is not a PNE. In fact, the following observation demonstrates that in many publishers' games induced by the PRP ranking function, no pure Nash equilibrium exists at all.

OBSERVATION 1. *Let $k \in \mathbb{N}$ and let G be a publishers' game with $n = 2$, $\lambda > \frac{1}{2}$, $d(x, y) = \frac{1}{k} \|x - y\|_2^2$, $r = r^*$, $x_0^1 = x_0^2 = \vec{0}_k$ and $x^* = \vec{1}_k$. Then, G possesses no PNE.*

PROOF. Assume by contradiction that $x^{eq} \in X$ is a PNE of G . We make use of the following auxiliary lemma (whose proof is given in Appendix E):

LEMMA 1. *Let G be a publishers' game with $r = r^*$ and $d(x, y) = \frac{1}{k} \|x - y\|_2^2$. If x^{eq} is a PNE of G , then $x_i^{eq} \in \{x^*, x_0^i\} \forall i \in N$.*

By Lemma 1, there are only 4 profiles that can be PNEs of G . We proceed to show that none of them is a PNE.

- $x^{eq} = (x^*, x_0^2)$: player 1 can improve her utility by deviating to $\hat{x}_1 = (\frac{1}{2}, \dots, \frac{1}{2})$.
- $x^{eq} = (x_0^1, x^*)$: player 2 can improve her utility by deviating to $\hat{x}_2 = (\frac{1}{2}, \dots, \frac{1}{2})$.
- $x^{eq} = (x_0^1, x_0^2)$: player 1 can improve her utility by deviating to $\hat{x}_1 = (\alpha, \dots, \alpha)$ for $0 < \alpha < \sqrt{\frac{1}{2\lambda}}$ since $u_1(\hat{x}_1, x_{-1}^{eq}) = 1 - \lambda\alpha^2 > 1 - \lambda\frac{1}{2\lambda} = \frac{1}{2} = u_1(x^{eq})$.
- $x^{eq} = (x^*, x^*)$: player 1 can improve her utility by deviating to x_0^1 since $u_1(x_0^1, x_{-1}^{eq}) = 0 - 0 > \frac{1}{2} - \lambda = u_1(x^{eq})$. □

As we discuss throughout the paper, the presence of a PNE holds significant importance for search engine designers. We now turn to suggest two other ranking function families, which will be shown (either theoretically or empirically) to be stable in terms of PNE existence and ε -better response dynamics convergence.

5.2 The Softmax Ranking Functions

Note that the instability of the PRP arises from the discontinuity of the argmax function, which introduces discontinuities into the ranking itself.⁷ Hence, a natural way to construct a smooth version of the PRP would be to replace argmax with its well-known smooth variant: the softmax function. Using softmax as a replacement for argmax is a standard practice in machine learning, where it serves as a differentiable approximation to discrete choices. This approach was successfully applied to strategic classification [25], item representation learning [32], information design [44], and beyond. In our context, the softmax ranking function is naturally defined as follows.⁸

DEFINITION 9. The **softmax** ranking function with inverse temperature constant $\beta > 0$, denoted by \tilde{r}^β , is the ranking function defined as follows:

$$\tilde{r}_i^\beta(x) := \frac{e^{-\beta \cdot d^*(x_i)}}{\sum_{j=1}^n e^{-\beta \cdot d^*(x_j)}} \quad (7)$$

Note that as $\beta \rightarrow \infty$, $\tilde{r}^\beta \rightarrow r^*$. This further explains the intuition behind the usage of the softmax function in our context and shows that β is a means to control how closely \tilde{r}^β aligns with the PRP. The class of softmax ranking functions is a subset of a wider class of ranking functions, known as the class of proportional ranking functions [27]:

DEFINITION 10. We say a ranking function r is **proportional** if there exists a continuously differentiable function $g : [0, 1] \rightarrow \mathbb{R}_{++}$ such that

$$r_i(x) = \frac{g(d^*(x_i))}{\sum_{j=1}^n g(d^*(x_j))} \quad (8)$$

The proportionality condition is an intuitive way to construct a function whose image is contained in the simplex Δ^n . We say a ranking function r **induces** a potential game if any publishers' game with r is a potential game. The following theorem shows that, despite its intuitive appeal, the class of proportional ranking functions (including the softmax ranking functions) is fundamentally incompatible with the desirable structure of potential games.

THEOREM 2. Except for the uniform ranking function $r_i \equiv \frac{1}{n}$, there exists no proportional ranking function that induces a potential game.

PROOF. Assume by contradiction that there exists a proportional ranking function r , defined by $r_i(x) = \frac{g(d^*(x_i))}{\sum_{j=1}^n g(d^*(x_j))}$ for some continuously differentiable $g : [0, 1] \rightarrow \mathbb{R}_{++}$, which is both not the uniform ranking function ($r_i \equiv \frac{1}{n}$) and induces a potential game. We show that the fact that r induces a potential game implies that g is constant, which implies r is the uniform ranking function and completes the proof.

Since r induces a potential game, any publishers' game with r as its ranking function is a potential game. In particular, any publishers' game with r as its ranking function, $k = 1$ and d twice continuously differentiable is a potential game. Such games, since $g > 0$, satisfy the conditions of Theorem 4.5 in Monderer and Shapley [31]. Therefore:

$$\frac{\partial^2 u_i(x)}{\partial x_i \partial x_j} = \frac{\partial^2 u_j(x)}{\partial x_i \partial x_j} \quad \forall i, j \in [n], \forall x \in X \quad (9)$$

⁷It is actually the argmin that is used by the PRP, but $\operatorname{argmin}_{i \in N} d^*(x_i) = \operatorname{argmax}_{i \in N} \{-d^*(x_i)\}$.

⁸Interestingly, the publishers' game induced by the softmax ranking function generalizes the well-known rent-seeking competition of Hirshleifer [20].

Recall $u_i(x) := r_i(x) - \lambda d_i^0(x_i)$. Since $\frac{\partial^2 \lambda d_i^0(x_i)}{\partial x_i \partial x_j} = \frac{\partial^2 \lambda d_j^0(x_i)}{\partial x_i \partial x_j} = 0$, we get:

$$\frac{\partial^2 r_i(x)}{\partial x_i \partial x_j} = \frac{\partial^2 r_j(x)}{\partial x_i \partial x_j} \quad \forall i, j \in [n], \forall x \in X \quad (10)$$

Note that r_i can be written as:

$$r_i(x) = \frac{(g \circ d^*)(x_i)}{\sum_{j=1}^n (g \circ d^*)(x_j)} \quad (11)$$

We now make use of the following lemma (whose proof is given in Appendix E):

LEMMA 2. Let $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be of the form $f_i(x) = \frac{g(x_i)}{\sum_{j=1}^n g(x_j)}$ for a continuously differentiable function $g : [0, 1] \rightarrow \mathbb{R}$.

Then

$$\frac{\partial^2 f_i(x)}{\partial x_i \partial x_j} = \frac{\partial^2 f_j(x)}{\partial x_i \partial x_j} \quad \forall i, j \in [n], \forall x \in [0, 1]^n \quad (12)$$

iff g is constant on $[0, 1]$.

By Lemma 2, $g \circ d^*$ is constant on $[0, 1]$. And since this is true for every twice continuously differentiable d^* , it is true in particular for $x^* = 0$, $d(a, b) = (a - b)^2$, which gives us $d^*(a) = a^2$. Note that in this case, the image of d^* is precisely $[0, 1]$. Therefore, g is constant on $[0, 1]$. □

Theorem 2 implies, in particular, that the softmax ranking function does not induce a potential game. One might naturally question whether the convergence of better response dynamics under the softmax ranking function is still guaranteed, despite the game not being a potential game (for example, there might exist a generalized ordinal potential). However, Observation 2, which uses the running example, demonstrates that this is not the case. While Observation 2 demonstrates this for a specific value of β , we hypothesize that for any given β , one can construct a game instance for which better-response dynamics diverge under the softmax ranking function with inverse temperature β .

OBSERVATION 2. There exist $\beta > 0$, a publishers' game instance G with $r = \tilde{r}^\beta$ and $\varepsilon > 0$ such that ε -better response dynamics may not converge.

PROOF. Consider the instance from Example 1, with the softmax ranking function with $\beta = 5$. Now, consider the following four strategy profiles (i 'th row is publisher i 's strategy):

$$x^{(1)} = \begin{pmatrix} 0.62959224 & 0.5189262 \\ 0.61146602 & 0.8503398 \end{pmatrix}, \quad x^{(2)} = \begin{pmatrix} 0.62959224 & 0.5033012 \\ 0.61146602 & 0.8503398 \end{pmatrix} \quad (13)$$

$$x^{(3)} = \begin{pmatrix} 0.62959224 & 0.5033012 \\ 0.62251457 & 0.83929126 \end{pmatrix}, \quad x^{(4)} = \begin{pmatrix} 0.62959224 & 0.5189262 \\ 0.62251457 & 0.83929126 \end{pmatrix} \quad (14)$$

Notice that $x^{(2)}$ is a unilateral deviation of Publisher 1 from $x^{(1)}$, $x^{(3)}$ is a unilateral deviation of Publisher 2 from $x^{(2)}$, $x^{(4)}$ is a unilateral deviation of Publisher 1 from $x^{(3)}$, and $x^{(1)}$ is a unilateral deviation of Publisher 2 from $x^{(4)}$. It is straightforward to show that each such deviation strictly improves the utility of the deviating publisher, and therefore, for a sufficiently small $\varepsilon > 0$, the four strategy profiles define a cyclic (hence not converging) ε -better response dynamics. □

Such instability naturally becomes more likely as β increases, causing the game to resemble the one induced by the PRP ranking function. However, as demonstrated in our empirical evaluation, there exists a range of β values where the softmax ranking function exhibits a remarkable degree of stability across nearly all game instances. Furthermore, for certain β values, the user welfare in the reached equilibrium surpasses the average user welfare achieved under the PRP ranking. This implies that, from a purely practical perspective, a designer can fine-tune the selection of β to achieve both user welfare and stability. Nevertheless, the lack of a theoretical guarantee regarding learning dynamics convergence motivates the study of linear ranking functions, presented hereinafter.

5.3 The Linear Ranking Functions

We begin by defining the linear relative relevance, which is the difference between the average distance of all other publishers from x^* and the distance of publisher i from x^* .

DEFINITION 11. Let $x \in X$ be a strategy profile in a publishers' game. The **linear relative relevance** of document (or publisher) $i \in N$ with respect to x is:

$$v_i(x) := \frac{1}{n-1} \left(\sum_{j \in N \setminus \{i\}} d^*(x_j) \right) - d^*(x_i) \quad (15)$$

Note that player i 's linear relative relevance increases as she becomes more relevant to the information need (minimizes $d^*(x_i)$) and decreases as any other publisher does so. A natural way to construct a ranking function using v would be expressions like $\hat{r}_i^{a,b}(x) := a \cdot v_i(x) + b$. Importantly, $\sum_{i=1}^n v_i \equiv 0$. Hence b must equal $\frac{1}{n}$ for the image of $\hat{r}^{a,b}$ to be contained in Δ^n . Further note that $v_i(x)$ might get the value -1 so we must have $a \leq \frac{1}{n}$ to ensure $a \cdot v_i(x) + \frac{1}{n} \geq 0$. Lastly, since a publisher with high v_i should get more exposure, $a \leq 0$ would not be plausible. This gives rise to the following definition:

DEFINITION 12. The **linear ranking function** with slope $a \in (0, \frac{1}{n}]$, denoted by \hat{r}^a , is the ranking function defined as follows:

$$\hat{r}_i^a(x) := a \cdot v_i(x) + \frac{1}{n} \quad (16)$$

A shortcoming of the linear ranking functions is evident from the above definition – the parameter a needs to be quite small as n goes large. This is problematic because it renders the value of r in a very small range, restricting the extent to which the ranking function incentivizes publishers to produce relevant content. However, as we highlighted when presenting the model, the case of small n is well-motivated, for example, in the case of niche domains. In such scenarios, the condition $a \in (0, \frac{1}{n}]$ is not particularly restrictive.⁹

We now proceed to provide some results about the significant advantages of the linear ranking function. Importantly, all linear ranking functions induce a potential game:

THEOREM 3. Let $a \in (0, \frac{1}{n}]$. \hat{r}^a induces a potential game, and the following is a potential function:

$$\hat{\phi}(x) = - \sum_{i=1}^n \left(a \cdot d^*(x_i) + \frac{1}{n} + \lambda d_i^0(x_i) \right) \quad (17)$$

PROOF. Let $g(t) = at + \frac{1}{n}$, and notice that the linear ranking function can now be written as $\hat{r}_i(x) = g(v_i(x))$. We will show that the following function is a potential function for the game G :

$$\phi(x) = - \sum_{i=1}^n \left(-g(d^*(x_i)) + \lambda d_i^0(x_i) \right) \quad (18)$$

⁹Additionally, in a real-life scenario, such a ranking function can incorporate an initial filtering of the documents based on a certain distance threshold, and only afterward apply the linear ranking function to determine the probability of the remaining documents being ranked first.

To show this, one should show that for any $i \in N$, $x_{-i} \in X_{-i}$, $x'_i, x''_i \in X_i$:

$$\phi(x'_i, x_{-i}) - \phi(x''_i, x_{-i}) = u_i(x'_i, x_{-i}) - u_i(x''_i, x_{-i}) \quad (19)$$

Let $i \in N$, $x_{-i} \in X_{-i}$, $x'_i, x''_i \in X_i$. Starting with the LHS, we get:

$$\begin{aligned} \phi(x'_i, x_{-i}) - \phi(x''_i, x_{-i}) &= -g(d^*(x'_i)) - \lambda d_i^0(x'_i) + g(d^*(x''_i)) + \lambda d_i^0(x''_i) \\ &= g(d^*(x''_i) - d^*(x'_i)) + \lambda(d_i^0(x''_i) - d_i^0(x'_i)) - \frac{1}{n} \end{aligned} \quad (20)$$

On the other hand, simplifying the RHS yields:

$$\begin{aligned} u_i(x'_i, x_{-i}) - u_i(x''_i, x_{-i}) &= g(v_i(x'_i, x_{-i})) - \lambda d_i^0(x'_i) - g(v_i(x''_i, x_{-i})) + \lambda d_i^0(x''_i) \\ &= g(v_i(x'_i, x_{-i}) - v_i(x''_i, x_{-i})) + \lambda(d_i^0(x''_i) - d_i^0(x'_i)) - \frac{1}{n} \\ &= g(d^*(x''_i) - d^*(x'_i)) + \lambda(d_i^0(x''_i) - d_i^0(x'_i)) - \frac{1}{n} \end{aligned} \quad (21)$$

The required equation (19) now follows directly from (20) and (21), which concludes the proof. \square

The following result follows from Theorem 1 and from the continuity of $\hat{\phi}$:

COROLLARY 1. *Let $a \in (0, \frac{1}{n}]$. Any publishers' game with $r = \hat{r}^a$ has at least one PNE. Moreover, for any $\varepsilon > 0$, any ε -better response dynamics converges.*

This result has major implications for the stability of the ecosystem. Specifically, if any ε -better response dynamics converges, then, in particular, any real-life better response dynamics converges. The result regarding learning dynamics convergence does not require full information of the initial documents of all players, but only assumes each player knows their initial document and the information need.

Note that whenever d is strictly biconvex, that is, $d(x, y)$ is strictly convex in x for any fixed y and vice versa (which is the case for semi-metrics like squared Euclidean norm), $\hat{\phi}$ is strictly concave. Together with the unique structure of the linear ranking functions, this allows the derivation of the following:

THEOREM 4. *Let $a \in (0, \frac{1}{n}]$. If d is strictly biconvex, then any publishers' game with $r = \hat{r}^a$ has a unique PNE and the equilibrium strategies are strictly dominant strategies.*

PROOF. Fix some publisher i and let us show she has a strictly dominant strategy. Player i 's utility is given by

$$\begin{aligned} u_i(x) &= r_i(x) - \lambda d_i^0(x_i) = \frac{v_i(x)}{n} + \frac{1}{n} - \lambda d_i^0(x_i) \\ &= \left(\sum_{j \in N \setminus \{i\}} \frac{d^*(x_j)}{n(n-1)} \right) - \frac{d^*(x_i)}{n} + \frac{1}{n} - \lambda d_i^0(x_i) \end{aligned} \quad (22)$$

Now, $\forall x_{-i} \in X_{-i}$:

$$\operatorname{argmax}_{x_i \in X_i} \{u_i(x_i, x_{-i})\} = \operatorname{argmax}_{x_i \in X_i} \left\{ -\frac{d^*(x_i)}{n} - \lambda d_i^0(x_i) \right\} \quad (23)$$

meaning that $\operatorname{argmax}_{x_i \in X_i} \{u_i(x_i, x_{-i})\}$ does not depend on x_{-i} . Note that since $g(x) = d(x, y)$ is strictly convex for any $y \in [0, 1]^k$, we get in particular that d_i^0 and d^* are strictly convex. This implies that $-\frac{d^*(x_i)}{n} - \lambda d_i^0(x_i)$ is strictly concave in x_i , and therefore has a unique maximizer. So we got that $\operatorname{argmax}_{x_i \in X_i} \{u_i(x_i, x_{-i})\}$ is a singleton that does not depend on x_{-i} . Let us denote $\operatorname{argmax}_{x_i \in X_i} \{u_i(x_i, x_{-i})\} = \{x_i^{dom}\}$. Then, the strategy x_i^{dom} is a strictly dominant strategy for player i by definition.

We proved that each player has a strictly dominant strategy. Note that the strategy profile where each player plays her strictly dominant strategy is a PNE. The uniqueness of the PNE follows from the fact that a strictly dominated strategy can never be an equilibrium strategy, and each player has only one strategy that is not strictly dominated (her dominant strategy). \square

The theorem shows that whenever d is strictly biconvex, not only do the linear ranking functions possess the strong behavioral property that any ε -better response learning dynamics converges, but also the resulting strategies are strictly dominant strategies for the publishers.

From now on, we focus our theoretical analysis on $\hat{r} := \hat{r}^{\frac{1}{n}}$. The choice $a = \frac{1}{n}$ is natural, as increasing the slope incentivizes the creation of content relevant to users, but we can use the same analysis method for other slopes. The following Lemma characterizes the unique equilibrium point in the case of the squared Euclidean semi-metric, in which the equilibrium strategy of each publisher is a weighted average of her initial document and x^* .¹⁰

LEMMA 3. *Any publishers' game with $r = \hat{r}$ and $d(x, y) = \frac{1}{k}\|x - y\|_2^2$ has a unique PNE, x^{eq} , given by $x_i^{eq} = \frac{x^* + \lambda n x_0^i}{1 + \lambda n}$. Moreover, x_i^{eq} is a strictly dominant strategy for any publisher $i \in N$.¹¹*

PROOF. Let G be a publishers' game with $r = \hat{r}$ and $d(x, y) = \frac{1}{k}\|x - y\|_2^2$. From Theorem 3, a potential function for G is:

$$\hat{\phi}(x) = - \sum_{i=1}^n \left(\frac{d^*(x_i)}{n} + \frac{1}{n} + \lambda d_i^0(x_i) \right) = - \sum_{i=1}^n \left(\frac{\|x_i - x^*\|_2^2}{nk} + \frac{1}{n} + \lambda \frac{\|x_i - x_0^i\|_2^2}{k} \right) \quad (24)$$

We notice that G is smooth and $\hat{\phi}$ is concave, therefore, a strategy profile is a PNE if and only if it is a potential maximizer (see, for example, [34]). In order to find a potential maximizer, we will find stationary points by setting the gradient of the potential with respect to x_i to 0_k . We will denote those points by x^{eq} .

$$\nabla_{x_i} \hat{\phi}(x) = \frac{-2}{nk} (x_i - x^*) - \frac{2\lambda}{k} (x_i - x_0^i) \quad (25)$$

$$\begin{aligned} 0_k = \nabla_{x_i} \hat{\phi}(x^{eq}) &\implies 0_k = -(x_i^{eq} - x^*) - \lambda n (x_i^{eq} - x_0^i) \\ &\implies 0_k = -\lambda n x_i^{eq} - x_i^{eq} + x^* + \lambda n x_0^i \\ &\implies \forall i \in N, x_i^{eq} = \frac{x^* + \lambda n x_0^i}{1 + \lambda n} \end{aligned} \quad (26)$$

Since $\hat{\phi}$ is strictly concave, x^{eq} is its unique maximizer and hence the only PNE in G . Note that the squared Euclidean norm is strictly biconvex.¹² Therefore by Theorem 4, x_i^{eq} is a strictly dominant strategy for every player i . \square

EXAMPLE 3. *Returning to our Example 1, Lemma 3 can be used to find the unique PNE in the game induced by the linear ranking function. In this example, we use the maximal slope $a = \frac{1}{n} = \frac{1}{2}$. The unique PNE of the game, in which the strategy of each player is also a dominant strategy, is given by:*

¹⁰In Appendix A we provide equilibrium analysis with L_2 norm.

¹¹Our model can be extended to an incomplete information setting where each publisher i is aware only of x^* and of x_0^i . One can show that x^{eq} is an equilibrium also in such games (i.e., an ex-post equilibrium).

¹²That is, $g(x) = \frac{1}{k}\|x - y\|_2^2$ is strictly convex for any $y \in [0, 1]^k$.

$$x_1^{eq} = \frac{x^* + \lambda n x_0^1}{1 + \lambda n} = \frac{(0.358, 0.908) + 1 \cdot 2 \cdot (0.950, 0.035)}{1 + 1 \cdot 2} \approx (0.753, 0.326) \quad (27)$$

$$x_2^{eq} = \frac{x^* + \lambda n x_0^2}{1 + \lambda n} = \frac{(0.358, 0.908) + 1 \cdot 2 \cdot (0.933, 0.773)}{1 + 1 \cdot 2} \approx (0.741, 0.818) \quad (28)$$

Notice that by Corollary 1 (and by the uniqueness of the equilibrium), any better-response dynamics will converge to this profile. Under this equilibrium, the linear relative relevance of Publisher 1 is given by:

$$\begin{aligned} v_1(x^{eq}) &= \frac{1}{2-1} d^*(x_2^{eq}) - d^*(x_1^{eq}) \\ &\approx \frac{1}{2} \|(0.741, 0.818) - (0.358, 0.908)\|^2 - \frac{1}{2} \|(0.753, 0.326) - (0.358, 0.908)\|^2 \\ &\approx 0.077 - 0.247 = -0.17 \end{aligned} \quad (29)$$

Correspondingly, the linear relative relevance of Publisher 2 is $v_2(x^{eq}) = -v_1(x^{eq}) \approx 0.017$. The utilities of the publishers (in equilibrium) are then given by:

$$\begin{aligned} u_1(x^{eq}) &= r_1(x^{eq}) - \lambda d_1^0(x_1^{eq}) \\ &\approx \frac{1}{2} \cdot v_1(x^{eq}) + \frac{1}{2} - 1 \cdot \frac{1}{2} \|(0.753, 0.326) - (0.950, 0.035)\|^2 \\ &\approx 0.415 - 0.062 = 0.353 \end{aligned} \quad (30)$$

$$\begin{aligned} u_2(x^{eq}) &= r_2(x^{eq}) - \lambda d_2^0(x_2^{eq}) \\ &\approx \frac{1}{2} \cdot v_2(x^{eq}) + \frac{1}{2} - 1 \cdot \frac{1}{2} \|(0.741, 0.818) - (0.933, 0.773)\|^2 \\ &\approx 0.585 - 0.019 = 0.566 \end{aligned} \quad (31)$$

The closed form of the PNE enables us to perform *welfare analysis*. For example, deriving the publishers' welfare in equilibrium by λ reveals that it has a global minimum in $\lambda = \frac{1}{n}$, that is, the point where the weight of x^* and x_0^i in the equilibrium strategy of player i is equal. It decreases in λ for $\lambda \leq \frac{1}{n}$ and increases for $\lambda \geq \frac{1}{n}$. As for the users' welfare, under a regularity condition on x_0 and x^* , it decreases in λ .¹³ Rigorous analysis is available within Appendix B.

We conclude this section by establishing that while applying a nonlinear function to the linear relative relevance might seem to generalize linear ranking functions, it turns out that this cannot be achieved.

THEOREM 5. *The only ranking functions of the form $r_i(x) = g(v_i(x))$ for some differentiable $g : [-1, 1] \rightarrow \mathbb{R}$ are for linear g .*

PROOF. If r is a valid ranking function, then in particular its image is contained in the simplex Δ^n for every publishers' game with $n = 3, k = 1, x^* = 0, d(a, b) = (a - b)^2$. Thus for every such game, $\forall x \in X = [0, 1]^3$ it holds that:

$$1 = \sum_{j=1}^n r_j(x) = \sum_{j=1}^n g(v_j(x)) \quad (32)$$

¹³The regularity condition depends on the values $d^i := d(x_0^i, x^*)$ for each $i \in N$, and its fulfillment is determined by the ratio between the mean and the variance of these distances. In the case of independently and uniformly distributed x_0 and x^* , the regularity condition holds with a very high probability. When the regularity condition does not hold, the users' welfare has a unique extremum, which is a global minimum point.

Define a function $f : X \rightarrow \mathbb{R}$ by $f(x) = \sum_{j=1}^3 g(v_j(x))$. Since $f \equiv 1$ on X , $\nabla f(x) = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}$ for all $x \in X$. Note that in such games, $d^*(a) = (a - 0)^2 = a^2$ and thus $d^{*'}(a) = 2a$. So $\forall x \in X, \forall i \in [n]$ it holds:

$$0 = \frac{\partial f(x)}{\partial x_i} = \sum_{j \neq i} g'(v_j(x)) \cdot \frac{2x_i}{n-1} - g'(v_i(x)) \cdot 2x_i \quad (33)$$

Denote by $\psi : \mathbb{R}^n \rightarrow \mathbb{R}^n$ the function defined by

$$\psi_i(d) := \frac{1}{n-1} \sum_{j \neq i} d_j - d_i \quad (34)$$

By (33), $\forall i \in N$:

$$0 = 2x_i \cdot \psi_i \left(\left(g'(v_l(x)) \right)_{l=1}^n \right) \quad (35)$$

Where $\left(g'(v_l(x)) \right)_{l=1}^n$ stands for the vector $(g'(v_1(x)) \quad g'(v_2(x)) \quad \dots \quad g'(v_n(x))) \in \mathbb{R}^n$.

Hence $\forall x \in X, \forall i, j \in [n]$ the following holds:

$$x_i \neq 0 \wedge x_j \neq 0 \implies \psi_i \left(\left(g'(v_l(x)) \right)_{l=1}^n \right) = 0 = \psi_j \left(\left(g'(v_l(x)) \right)_{l=1}^n \right) \quad (36)$$

We now use the following lemma (whose proof appears in Appendix E):

LEMMA 4. Let $\psi : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be the function defined by $\psi_i(d) = \frac{1}{n-1} \sum_{j \neq i} d_j - d_i$.

Then $\forall d \in \mathbb{R}^n, \forall i, j \in [n]$,

$$\psi_i(d) = \psi_j(d) \iff d_i = d_j \quad (37)$$

By Lemma 4, $\forall x \in X, \forall i, j \in [n]$ it holds that:

$$x_i \neq 0 \wedge x_j \neq 0 \implies g'(v_i(x)) = g'(v_j(x)) \quad (38)$$

Now, we prove that g must be linear by proving that g' must be constant on $[-1, 1]$. We do this in two steps: first, we show that g' is constant on $(-0.75, 0.75)$ and then we show it is constant on the entire $[-1, 1]$ as well.

For the first step, we look at the profiles $x_\varepsilon := (\sqrt{0.5 - \varepsilon}, \sqrt{0.5}, \sqrt{0.5 + \varepsilon})$ for $\varepsilon \in (0, 0.5)$. The linear relative relevance vector in this case is $v(x_\varepsilon) = (1.5\varepsilon, 0, -1.5\varepsilon)$. By applying (38) on $x_\varepsilon, i = 1, j = 2$ we get $g'(1.5\varepsilon) = g'(0)$. By applying (38) on $x_\varepsilon, i = 2, j = 3$ we get $g'(-1.5\varepsilon) = g'(0)$. So, we got $\forall \varepsilon \in (0, 0.5), g'(-1.5\varepsilon) = g'(0) = g'(1.5\varepsilon)$. This means that $\forall t \in (0, 0.75), g'(-t) = g'(0) = g'(t)$. Therefore g' is constant on $(-0.75, 0.75)$.

Now for the second step, let $t \in (-1, 1)$ and we show that $g'(t) = g'(0)$. Consider the profile $\hat{x} := \left(\sqrt{\frac{1-t}{2}}, \sqrt{\frac{1+t}{2}}, \sqrt{\frac{1+t}{2}} \right)$.

The linear relative relevance vector in this case is $v(\hat{x}) = (t, -\frac{t}{2}, -\frac{t}{2})$. By applying (38) on $\hat{x}, i = 1, j = 2$ we get $g'(t) = g'(-\frac{t}{2})$. Since $t \in (0, 1), -\frac{t}{2} \in (-0.75, 0.75)$, so by the first step we have that $g'(-\frac{t}{2}) = g(0)$. So $g'(t) = g'(0)$ as we wanted to show. This proves that g' is constant on $(-1, 1)$. Therefore, g is linear on $(-1, 1)$. Since g is continuous, it is also linear on $[-1, 1]$. □

6 Empirical Results

In this section, we evaluate the performance of the PRP, the linear, and the softmax ranking functions in simulated environments that aim to reflect realistic better response dynamics. Empirical analysis of competitive retrieval settings is a difficult challenge, as discussed in Kurland and Tennenholtz [24]. Conducting experiments using real data requires the implementation of each proposed ranking function into real systems, in which publishers repeatedly compete for exposure and receive feedback from a system that ranks according to the proposed approach. Therefore, we evaluate our ranking functions using a code-simulated environment which we call the **Discrete Better Response Dynamics Simulation**.¹⁴ These dynamics are typical in real-world scenarios, where publishers endeavor to enhance the visibility and impressions of their content within the platform they operate. We focus on the canonical of $n = 2$, meaning that two prominent publishers compete for search engine impressions. See Appendix D.1 for experiments and discussion beyond the two-publishers scenario.

6.1 Simulation Details

In the Discrete Better Response Dynamics Simulation, publishers start with their initial documents, and at each timestep, one publisher modifies their document to improve their utility. The modification is done by taking a step that maximizes the publisher's utility, from a predefined subset of possible directions \mathcal{D} and step sizes \mathcal{S} . Formally, in each simulation, we consider a publishers' game G with $n = 2$, $k = 2$,¹⁵ $\lambda \in [0, 1]$, $d(x, y) = \frac{1}{k} \|x - y\|_2^2$ and uniformly drawn¹⁶ $x^* \in [0, 1]^k$ and $x_0 \in [0, 1]^{n \cdot k}$. A Discrete Better Response Dynamics Simulation is said to **converge** if it has reached an ε -PNE under the discreteness constraints (the constraints of predefined directions and step sizes) within no more than T iterations.

To evaluate the performance of the different ranking functions, we have estimated the *convergence ratio*—the proportion of simulations that converged—and the expected long-term publishers' and users' welfare. In case the simulation converged, we estimate the welfare measures over the long term using the welfare measures of the profile the simulation converged to. Otherwise, we average the welfare measures over the last M rounds of the simulation, where M is an additional simulation parameter. A justification for this estimation method is the empirically observed periodic nature of the dynamics that did not converge (see §6.4 for a detailed discussion). For each value of λ , we performed 500 simulations and constructed a bootstrap confidence interval with a confidence level of 95%. Algorithm 1 provides the pseudo-code for a single simulation.¹⁷

6.2 Comparison of the PRP, Softmax, and Linear Ranking Functions

Figure 1 provides a comparative analysis of the three ranking functions: PRP, softmax (with $\beta = 1$), and linear (with $a = \frac{1}{n}$). The rightmost chart in Figure 1 highlights a fundamental shortcoming of the PRP ranking - most of the dynamics it induces do not converge. However, dynamics under the other two ranking functions always converge. This is a key advantage of these functions. A noteworthy point that emerged from the simulations but is not shown in the graphs is that convergence to an ε -PNE under the softmax and linear ranking functions was very rapid, typically occurring after only a few timesteps.

Another insight from Figure 1 is that the users' welfare of the PRP is roughly constant across λ values. A plausible explanation is that in the case of PRP and the tested λ range, λ has little, if any, impact on the behavior of publishers. This conjecture might also explain why the publishers' welfare of the PRP appears to linearly

¹⁴Code and results are available in <https://github.com/ireinman/The-Search-for-Stability>.

¹⁵The time complexity of each simulation timestep grows exponentially with k , therefore we used $k = 2$. An analysis of the effect of k on the results can be found within Appendix C.

¹⁶Experiments with other distributions of x^* and x_0 are reported in Appendix D.2.

¹⁷We used $\mathcal{S} = \{2^{-6}, 2^{-5}, \dots, 2^{-2}, 0.5, 0.6, \dots, 1\}$, $\mathcal{D} = \left\{ \frac{d}{\|d\|_2} : d \in \{-1, 0, 1\}^k \setminus \{\vec{0}\} \right\}$, $T = 1000$, $M = 900$, $\varepsilon = 10^{-6}$, and a bootstrap sample size of $B = 500$.

Algorithm 1 Discrete Better Response Dynamics Simulation

```

1: Input:  $G, \mathcal{S}, \mathcal{D}, T, M, \epsilon$ 
2:  $x^{(0)} \leftarrow x_0$ 
3:  $conv \leftarrow False$ 
4: for  $t = 1, 2, \dots, T$  do
5:    $\mathcal{I} \leftarrow$  the set of players that can improve their utility
6:   by more than  $\epsilon$  by taking a step in direction  $d \in \mathcal{D}$ 
7:   with step size  $s \in \mathcal{S}$ 
8:   if  $\mathcal{I} \neq \emptyset$  then
9:     Draw  $i \in \mathcal{I}$  uniformly
10:     $BR_i \leftarrow \operatorname{argmax}_{(s,d) \in \mathcal{S} \times \mathcal{D}} \{u_i(x_i^{(t-1)} + s \cdot d, x_{-i}^{(t-1)})\}$ 
11:    Draw  $(s', d') \in BR_i$  uniformly
12:     $x^{(t)} \leftarrow (x_i^{(t-1)} + s' \cdot d', x_{-i}^{(t-1)})$ 
13:   else
14:      $conv \leftarrow True$ 
15:   return  $(\mathcal{U}(x^{(t-1)}), \mathcal{V}(x^{(t-1)}), conv)$ 
16:   end if
17: end for
18:  $T_s \leftarrow T - M + 1$ 
19: return  $(\operatorname{Avg}\{\mathcal{U}(x^{(t)})\}_{t=T_s}^T, \operatorname{Avg}\{\mathcal{V}(x^{(t)})\}_{t=T_s}^T, conv)$ 

```

decrease with λ : the dynamics remain the same, but the term $-\lambda \cdot \sum_{i=1}^n d_i^0(x_i)$ in the publishers' welfare gains more influence as λ increases. Note that the PRP ranking function appears to sacrifice stability in pursuit of increased users' welfare.

In the linear ranking function, the effect of λ on the welfare measures matches the theoretical results presented in §5.3—the publishers' welfare obtains a minimum in $\lambda = \frac{1}{n}$, and the users' welfare is monotonically decreasing in λ .¹⁸ Another observation from Figure 1 is that the softmax ranking function exhibits similar trends. In this figure, an inherent trade-off in our model can be seen - an increase in publishers' welfare leads to a decrease in users' welfare and vice versa.

It seems that the softmax ranking function has higher publishers' welfare and lower users' welfare than the linear ranking function. Yet, tuning the inverse temperature constant β or the slope a can change this. We now provide an empirical analysis of these hyperparameters' influence.

6.3 The Effect of Hyperparameters Selection

Figure 2 provides a comparative analysis of the influence of the inverse temperature constant β in the softmax ranking function. It is evident that β serves as a means to control the inherent trade-off between publishers' welfare and users' welfare. A search engine designer can tune β according to the relative importance of the two welfares within her domain and according to her estimation of λ . Importantly, the convergence ratio of \tilde{r}^β is 1 for all tested values of β . Therefore, the softmax with high β values (for example, $\beta \in \{5, 10\}$) manages to surpass

¹⁸The likelihood of the regularity condition mentioned in §5.3 under uniform distribution of x_0 and x^* is very close to 1. This explains why the results shown in Figure 1 correlate with the theoretical result for when the condition does hold.

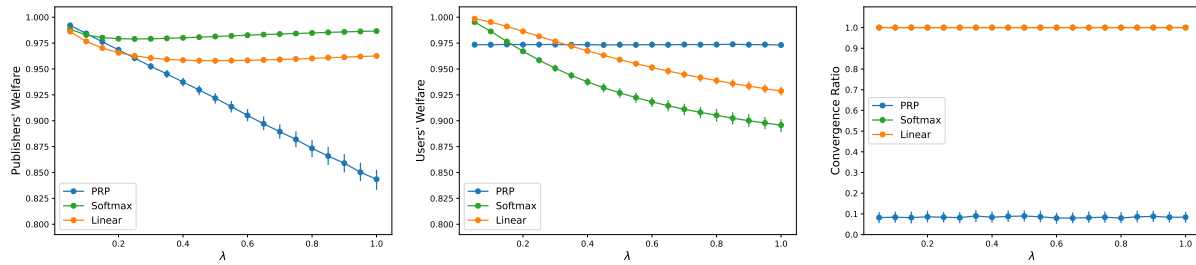


Fig. 1. Comparison of the PRP, softmax ($\beta = 1$) and linear ($a = \frac{1}{n} = 0.5$) ranking functions. The linear and the softmax ranking functions had a convergence ratio of 1.

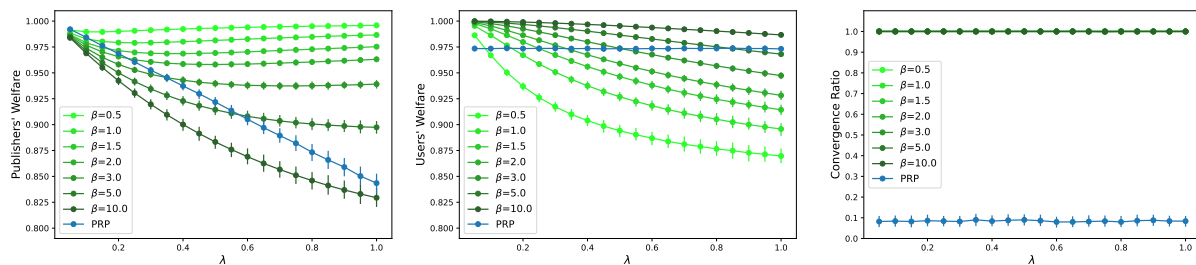


Fig. 2. The effect of β on the welfare measures and the convergence ratio in the softmax ranking function. All softmax ranking functions had a convergence ratio of 1.

the PRP in terms of users' welfare while maintaining the desired property of 100% convergence. The cost is sacrificing publishers' welfare. Therefore, we claim that the softmax ranking functions offer viable alternatives to the PRP function.

As mentioned in §5, the PRP ranking function maximizes the users' welfare for a fixed profile. The reason why softmax ranking functions with high β values nevertheless manage to achieve higher users' welfare than the PRP is that the PRP is only short-term optimal, while the softmax ranking functions preserve high users' welfare in the long term as well. Unlike the softmax ranking functions, most of the dynamics the PRP induces do not converge. In such ecosystems, the publishers repeatedly reach profiles with sub-optimal users' welfare, and this harms the long-term users' welfare of the PRP. This emphasizes the need for the search for stability.

Now, let us delve deeper into the linear ranking function and compare it to the softmax. Figure 3 provides a comparative analysis of the influence of the slope a in the linear ranking function. As was the case for the softmax ranking function, \hat{r}^a has a convergence ratio of 1 for all tested a values. This is not surprising, as we proved it in Corollary 1.

Just like β in the softmax functions, the slope a allows the search engine designer to control the publisher-user trade-off. There is, however, a core difference between the tuning of a and β . While the range of a values for which \hat{r}^a is a valid ranking function is bounded from above by $\frac{1}{n}$ (see §5.3), one can set β to be arbitrarily large. This means that using the softmax ranking function it is possible to achieve high users' welfare values that cannot be obtained using the linear ranking function. Nevertheless, the linear still has a significant advantage - it induces a potential game. As shown in §5.3, the existence of a potential function enables rigorous theoretical analysis and guarantees the convergence of ϵ -better response dynamics.

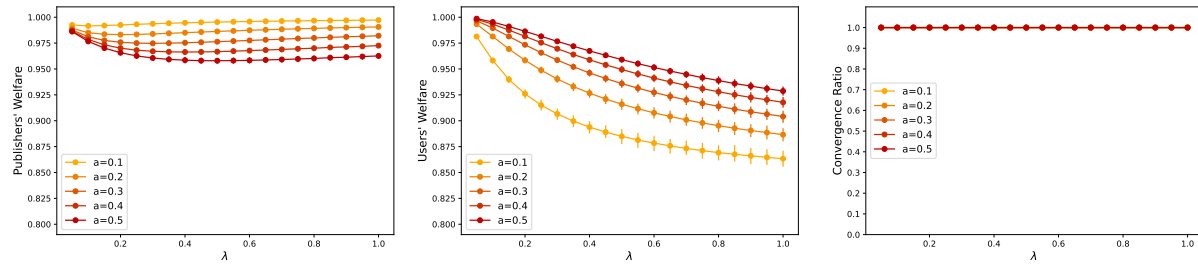


Fig. 3. The effect of a on the welfare measures and the convergence ratio in the linear ranking function. All linear ranking functions had a convergence ratio of 1.

6.4 Learning Dynamics Under the PRP Ranking Function

In this section, we delve into the learning dynamics of games induced by the PRP ranking function. While the dynamics in games induced by the softmax and the linear ranking functions are relatively easy to anticipate, as they converge to an ε -PNE in a few timesteps, PRP-based games often exhibit completely different phenomena. We will explore these intriguing dynamics in detail.

During the examination of the PRP ranking function, we noticed that the vast majority (about 99.5%) of the dynamics it induces belong to one of two types of dynamics.¹⁹ One type, which constitutes a minority of the cases, is when the simulation converges already in the first round because no publisher can improve their utility under the constraints of the discrete simulation. In those cases, the publisher whose initial document is farther from x^* (the information need) cannot get closer to x^* than the other publisher in a single step, and therefore cannot improve her utility in a single step. Moreover, the other publisher gets the maximum utility, 1, because she is the closest to x^* and she is at her initial document, therefore she also cannot improve her utility in a single step. In the other type of simulations, which constitutes most of the cases, the publishers enter a pseudo-periodic²⁰ behavior that never converges. Figure 4 presents the welfare measures during one representative instance of the latter type of simulations.

The periodic nature of the publishers' behavior is evident in Figure 4. In each cycle, there is a period where the publishers compete to provide the closest document to x^* . Each of them repeatedly changes her document (strategy) by the smallest margin sufficient for her to get closer to the information need than her competitor. This competition, which is accompanied by an increase in the users' welfare and a decrease in the publishers' welfare, continues until one of the players resigns and abruptly returns as close as she can to her initial document. Her decision to resign is either because the discrete nature of the simulation prevents her from bypassing the other publisher, or because further bypassing the other publisher is not beneficial for her since it positions her too far from her initial document. Immediately after this resignation, the publisher who won the competition takes a big step toward her initial document while still maintaining her position as the closest to x^* , since as long as she is still closer to x^* than the other publisher, getting closer to her initial document is trivially beneficial for her. This mutual withdrawal is accompanied by a sharp decrease in users' welfare and a sharp increase in publishers' welfare. When the losing publisher notices the winning publisher's retreat from x^* , she exploits the situation to get back into the competition and bypass her rival. The entire process is then repeated over and over.

¹⁹In the rare dynamics that do not fit into either of the two main types of dynamics, the simulation converges after more than one timestep. The convergence in those cases can be attributed to the fact that the step sizes are bounded from below.

²⁰We use the term "pseudo-periodic" instead of "periodic" since the publisher who deviates in each timestep is chosen at random from all the publishers who can improve their utility by more than ε . Because of this randomness, the dynamics that exhibit a periodic nature usually display variations between cycles, thus justifying the term "pseudo-periodic".

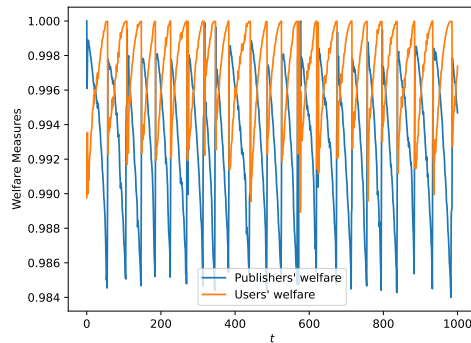


Fig. 4. The welfare measures during a Discrete Better Response Dynamics Simulation with the PRP ranking function on a representative instance of a publishers' game with $\lambda = 0.5$, $n = 2$ and $k = 2$. The figure illustrates the cyclic nature of the dynamics over time.

It can be seen that the PRP reaches profiles with users' welfare of almost 1, which is the optimal users' welfare value, but it does not stay there. This explains why, despite the PRP being optimized for short-term relevance (that is, for a fixed profile, it maximizes the users' welfare), its instability harms its long-term performance. This phenomenon once again emphasizes the need for the search for stability.

The periodicity is best observed when animating the documents in a two-dimensional embedding space, which can be done using the code provided in the supplementary material. An example video is provided as well.²¹ In this visualization, one can notice that during the simulation, each publisher stays close to the line segment that connects x^* with her initial document. This phenomenon is correlated with the theoretical result of Lemma 5 (in Appendix A), which determines that in games induced by various ranking functions, including the PRP, any strategy that is outside this line segment is strictly dominated.²²

7 Discussion

We studied a game-theoretic model for information retrieval, which models documents as vectors in a continuous space, accounts for original content, and distinguishes between publishers' and users' welfare. We have provided both theoretical and empirical negative results about the stability of ecosystems in which the PRP is used to rank the documents in the corpus and demonstrated how instability harms long-term users' welfare. Driven by the search for stability, we propose two alternative ranking function families - the linear and the softmax ranking function families. We focus on a strong criterion for stability, which is the convergence of better response dynamics. For the linear ranking function, we have theoretically established the convergence of ε -better response dynamics, and we have empirically demonstrated that the softmax ranking function fosters stable ecosystems. The convergence of any ε -better response dynamics in games induced by the linear ranking function is a critical finding, as such dynamics mirror the competitive landscape observed in the real world among publishers. Moreover, we empirically showed a trade-off between publishers' welfare and users' welfare and provided means for a search engine designer to control it.

²¹The video can be downloaded from our GitHub repository: https://github.com/ireinman/The-Search-for-Stability/blob/main/figures/animation_example.mp4.

²²The fact that during the simulation players do not play exactly on the line segment but rather just close to it is due to the discrete nature of our simulation.

Solution concept and stability guarantee. A natural question in the context of our analysis is regarding the consideration of a *pure* Nash equilibrium existence as a desirable property. One may claim that the existence of a *mixed* Nash equilibrium, in which publishers are allowed to engage in mixed strategies (that is, randomize pure actions), is sufficient as a stability criterion of the ecosystem induced by the ranking scheme. We provide several justifications for our rather stronger requirement.

First, while it may be plausible that publishers engage in mixed strategies (for instance, by using generative AI for content creation), a mixed Nash equilibrium results in unpredictable outcomes for the designer. In contrast, a PNE reflects a scenario where all publishers play deterministically, leading to a more stable ecosystem. Predicting and monitoring the content available within the platform becomes easier when publishers converge to a PNE. Second, we highlight that when publishers engage with pure strategies, users can be sure they will always find relevant content on the platform. Indeed, under some different modeling assumptions, generating content deterministically may be less desirable (for example, users may become bored if content is generated deterministically). In some contexts, this is not the case. For instance, whenever x^* is a piece of information that the user expects to find in a search engine, she always wants to find the most informative document, and boredom does not play a role. This stands in contrast to other use cases, such as content consumed for entertainment purposes. Another reason to prefer PNE is the effort required for generating randomized content. For instance, randomization can be done using LLMs, which is expensive when used at scale. See also the discussion in Ben-Porat and Tennenholtz [10] on PNE as a property indicating stability.

Importantly, much of our analysis focuses not only on the existence of a PNE but also on a stronger stability property: the guarantee that any better-response dynamics will converge. This property is particularly powerful in real-world applications, where the specific structure of content modifications—such as how changes in the text space translate to the embedding space—is often unknown to the designer. Anticipating these transformations is challenging, but under the mild assumption that publishers can modify documents to improve their utility (as determined by the embedding space), strategic content modifications are guaranteed to converge, even without the designer explicitly understanding their precise structure.

Limitations and future work. We find it important to discuss several limitations of this work. First, our model assumes a specific form of the publishers' utilities (winning probability minus integrity cost). Despite being restrictive, the assumption is quite standard, as the two terms usually do not interact. We assume publishers' full knowledge of the ranking function and the information need. This is a very common assumption when publishers compete in a *niche domain* – such publishers are usually experts who are familiar with the requirements and needs of users in their domain, as well as with the platform's algorithm. Following the literature on information retrieval games, our model assumes that publishers only gain utility from being ranked first. Additionally, we experimented with a specific better-response dynamics structure and specific distance functions. Importantly, for cases where any better response dynamics is guaranteed to converge to a unique equilibrium, these specifications are not restrictive at all. Another limitation in the experimental setting is the restriction of uniformly i.i.d. distributing x_0 and x^* . In Appendix D.2 we provide preliminary results beyond this case, in which all trends agree with the ones presented in the paper.

Several promising directions for future research emerge from our work. First, extending the model to incorporate heterogeneous cost parameters (that is, assigning each publisher an individual λ_i rather than assuming a common λ) would better capture the asymmetries found in real-world publisher behavior. We note that many of our results and techniques naturally extend to this more general setting.

Second, a more refined analysis of the *rate of convergence* to equilibrium under various ranking functions is of both theoretical and practical interest. In real-world systems, designers often care about the speed at which a stable state is reached. Furthermore, while some information needs are timeless, others are time-sensitive, making quick convergence essential for maintaining user relevance (for instance, consider news or trending

topics). Modeling such *time-dependent preferences* among publishers and users presents a valuable avenue for future work.

A related extension is to study *non-stationary information needs*, that is, scenarios in which the target query itself evolves over time, potentially as a function of previously published documents. This captures the well-documented phenomenon of *preference shifts* in recommendation systems [17]. Formally, such a setting could be modeled using the framework of stochastic games, and we leave its integration into the publishers' game model as an interesting open direction.

Finally, while our model reflects considerations such as *integrity toward a publisher's preferred ideal content* (as captured by distance from the initial document x_i^0), an alternative approach would define cost in terms of *modification effort*—that is, the distance from the previously published document. This alternative cost model may better reflect scenarios where the dominant constraint is the *effort or friction* involved in repeatedly editing or rewriting content.

Acknowledgments

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). We would like to thank the associate editor and the anonymous reviewers for their helpful comments. We also thank Yotam Gafni for his helpful comments on an early version of the paper. All authors contributed equally to this work.

References

- [1] M. Abbasi, A. Bhaskara, and S. Venkatasubramanian. 2021. Fair clustering via equitable group representations. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (FAccT '21). Association for Computing Machinery, Virtual Event, Canada, 504–514. ISBN: 9781450383097. DOI: [10.1145/3442188.3445913](https://doi.org/10.1145/3442188.3445913).
- [2] A. Agrawal, A. Ali, S. Boyd, et al. 2021. Minimum-distortion embedding. *Foundations and Trends® in Machine Learning*, 14, 3, 211–378. DOI: [10.1561/2200000090](https://doi.org/10.1561/2200000090).
- [3] S. Alshamrani. 2024. Distance matters: euclidean embedding distances for improved language model generalization and adaptability. *IEEE Access*, 12, 103583–103593. DOI: [10.1109/ACCESS.2024.3434612](https://doi.org/10.1109/ACCESS.2024.3434612).
- [4] G. Bahar, O. Ben-Porat, K. Leyton-Brown, and M. Tennenholtz. 2020. Fiduciary bandits. In *Proceedings of the 37th International Conference on Machine Learning* (Proceedings of Machine Learning Research). H. D. III and A. Singh, (Eds.) Vol. 119. PMLR, (July 2020), 518–527. <https://proceedings.mlr.press/v119/bahar20a.html>.
- [5] R. Ben Basat, M. Tennenholtz, and O. Kurland. 2017. A game theoretic analysis of the adversarial retrieval setting. *Journal of Artificial Intelligence Research*, 60, 1127–1164. DOI: <https://doi.org/10.1613/jair.5547>.
- [6] R. Ben Basat, M. Tennenholtz, and O. Kurland. 2015. The probability ranking principle is not optimal in adversarial retrieval settings. In *Proceedings of the 2015 International Conference on The Theory of Information Retrieval* (ICTIR '15). Association for Computing Machinery, Northampton, Massachusetts, USA, 51–60. ISBN: 9781450338332. DOI: [10.1145/2808194.2809456](https://doi.org/10.1145/2808194.2809456).
- [7] O. Ben-Porat, G. Goren, I. Rosenberg, and M. Tennenholtz. 2019. From recommendation systems to facility location games. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33, 01, (July 2019), 1772–1779. DOI: [10.1609/aaai.v33i01.33011772](https://doi.org/10.1609/aaai.v33i01.33011772).
- [8] O. Ben-Porat, I. Rosenberg, and M. Tennenholtz. 2019. Convergence of learning dynamics in information retrieval games. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33, 01, (July 2019), 1780–1787. DOI: [10.1609/aaai.v33i01.33011780](https://doi.org/10.1609/aaai.v33i01.33011780).
- [9] O. Ben-Porat, F. Sandomirskiy, and M. Tennenholtz. 2021. Protecting the protected group: circumventing harmful fairness. *Proceedings of the aaai conference on artificial intelligence*, 35, 6, (May 2021), 5176–5184. DOI: [10.1609/aaai.v35i6.16654](https://doi.org/10.1609/aaai.v35i6.16654).
- [10] O. Ben-Porat and M. Tennenholtz. 2018. A game-theoretic approach to recommendation systems with strategic content providers. *Advances in Neural Information Processing Systems*, 31. S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, (Eds.) https://proceedings.neurips.cc/paper_files/paper/2018/file/a9a1d5317a33ae8cef33961c34144f84-Paper.pdf.
- [11] D. P. Bertsekas. 1997. Nonlinear programming. *Journal of the Operational Research Society*, 48, 3, 334–334. eprint: <https://doi.org/10.1057/palgrave.jors.2600425>. DOI: [10.1057/palgrave.jors.2600425](https://doi.org/10.1057/palgrave.jors.2600425).
- [12] O. Butman, A. Shtok, O. Kurland, and D. Carmel. 2013. Query-performance prediction using minimal relevance feedback. In *Proceedings of the 2013 Conference on the Theory of Information Retrieval* (ICTIR '13). Association for Computing Machinery, Copenhagen, Denmark, 14–21. ISBN: 9781450321075. DOI: [10.1145/2499178.2499201](https://doi.org/10.1145/2499178.2499201).

- [13] M. Cary, A. Das, B. Edelman, I. Giotis, K. Heimerl, A. R. Karlin, S. D. Kominers, C. Mathieu, and M. Schwarz. 2014. Convergence of position auctions under myopic best-response dynamics. *ACM Transactions on Economics and Computation (TEAC)*, 2, 3, Article 9, (July 2014), 20 pages. DOI: [10.1145/2632226](https://doi.org/10.1145/2632226).
- [14] F. Chierichetti, R. Kumar, S. Lattanzi, and S. Vassilvitskii. 2017. Fair clustering through fairlets. In *Advances in Neural Information Processing Systems*. I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, (Eds.) Vol. 30. Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2017/file/978f5b0cc4eccc88ad48ce3914124a2-Paper.pdf.
- [15] G. G. Chowdhury. 2010. *Introduction to modern information retrieval*. Facet publishing.
- [16] C. Claus and C. Boutilier. 1998. The dynamics of reinforcement learning in cooperative multiagent systems. *AAAI/IAAI*, 1998, 746-752, 2.
- [17] S. Dean and J. Morgenstern. 2022. Preference dynamics under personalized recommendations. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC '22)*. Association for Computing Machinery, Boulder, CO, USA, 795–816. ISBN: 9781450391504. DOI: [10.1145/3490486.3538346](https://doi.org/10.1145/3490486.3538346).
- [18] Y. Freund and R. E. Schapire. 1999. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29, 1, 79–103. DOI: <https://doi.org/10.1006/game.1999.0738>.
- [19] M. Hardt, N. Megiddo, C. Papadimitriou, and M. Wootters. 2016. Strategic classification. In *Proceedings of the 2016 ACM conference on innovations in theoretical computer science (ITCS '16)*. Association for Computing Machinery, Cambridge, Massachusetts, USA, 111–122. ISBN: 9781450340571. DOI: [10.1145/2840728.2840730](https://doi.org/10.1145/2840728.2840730).
- [20] J. Hirshleifer. 1989. Conflict and rent-seeking success functions: ratio vs. difference models of relative success. *Public Choice*, 63, 2, 101–112. DOI: [10.1007/BF00153394](https://doi.org/10.1007/BF00153394).
- [21] J. Hron, K. Krauth, M. I. Jordan, N. Kilbertus, and S. Dean. 2023. Modeling content creator incentives on algorithm-curated platforms. (2023). <https://arxiv.org/abs/2206.13102> arXiv: 2206.13102 [cs.GT].
- [22] M. Jagadeesan, N. Garg, and J. Steinhardt. 2023. Supply-side equilibria in recommender systems. In *Advances in Neural Information Processing Systems*. A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, (Eds.) Vol. 36. Curran Associates, Inc., 14597–14608. https://proceedings.neurips.cc/paper_files/paper/2023/file/2f1486343c2c942a617e4f5bb0cc64c8-Paper-Conference.pdf.
- [23] T. Joachims, L. Granka, B. Pan, H. Hembrooke, and G. Gay. 2017. Accurately interpreting clickthrough data as implicit feedback. In *Acm Sigir Forum* number 1. Vol. 51. Acm New York, NY, USA, 4–11.
- [24] O. Kurland and M. Tennenholtz. 2022. Competitive search. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '22)*. Association for Computing Machinery, Madrid, Spain, 2838–2849. ISBN: 9781450387323. DOI: [10.1145/3477495.3532771](https://doi.org/10.1145/3477495.3532771).
- [25] S. Levanon and N. Rosenfeld. 2022. Generalized strategic classification and the case of aligned incentives. In *Proceedings of the 39th International Conference on Machine Learning (Proceedings of Machine Learning Research)*. K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, and S. Sabato, (Eds.) Vol. 162. PMLR, (July 2022), 12593–12618. <https://proceedings.mlr.press/v162/levanon22a.html>.
- [26] C. Liu and Y. Wei. 2016. The impacts of time constraint on users' search strategy during search process. *Proceedings of the Association for Information Science and Technology*, 53, 1, 1–9. DOI: [10.1002/pra2.2016.14505301051](https://doi.org/10.1002/pra2.2016.14505301051).
- [27] O. Madmon, I. Pipano, I. Reinman, and M. Tennenholtz. 2025. On the convergence of no-regret dynamics in information retrieval games with proportional ranking functions. In *The Thirteenth International Conference on Learning Representations*. <https://openreview.net/forum?id=jjXZvPe5z0>.
- [28] R. Meir, M. Polukarov, J. Rosenschein, and N. Jennings. 2010. Convergence to equilibria in plurality voting. *Proceedings of the AAAI conference on artificial intelligence*, 24, 1, (July 2010), 823–828. DOI: [10.1609/aaai.v24i1.7624](https://doi.org/10.1609/aaai.v24i1.7624).
- [29] B. Mitra and N. Craswell. 2017. Neural models for information retrieval. (2017). <https://arxiv.org/abs/1705.01509> arXiv: 1705.01509 [cs.IR].
- [30] M. Mladenov, E. Creager, O. Ben-Porat, K. Swersky, R. Zemel, and C. Boutilier. 2020. Optimizing long-term social welfare in recommender systems: a constrained matching approach. In *Proceedings of the 37th International Conference on Machine Learning (Proceedings of Machine Learning Research)*. H. D. III and A. Singh, (Eds.) Vol. 119. PMLR, (July 2020), 6987–6998. <https://proceedings.mlr.press/v119/mladenov20a.html>.
- [31] D. Monderer and L. S. Shapley. 1996. Potential games. *Games and Economic Behavior*, 14, 1, 124–143. DOI: <https://doi.org/10.1006/game.1996.0044>.
- [32] O. Nahum, G. Noti, D. Parkes, and N. Rosenfeld. 2024. Decongestion by representation: learning to improve economic welfare in marketplaces. (2024). <https://arxiv.org/abs/2306.10606> arXiv: 2306.10606 [cs.LG].
- [33] V. Nair, G. Ghalme, I. Talgam-Cohen, and N. Rosenfeld. 2022. Strategic representation. In *Proceedings of the 39th International Conference on Machine Learning (Proceedings of Machine Learning Research)*. K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, and S. Sabato, (Eds.) Vol. 162. PMLR, (July 2022), 16331–16352. <https://proceedings.mlr.press/v162/nair22a.html>.
- [34] A. Neyman. 1997. Correlated equilibrium and potential games. *International Journal of Game Theory*, 26, 2, 223–227. DOI: [10.1007/BF01295851](https://doi.org/10.1007/BF01295851).

- [35] K. Nissim, R. Smorodinsky, and M. Tennenholtz. 2018. Segmentation, incentives, and privacy. *Mathematics of Operations Research*, 43, 4, 1252–1268. doi: [10.1287/moor.2017.0903](https://doi.org/10.1287/moor.2017.0903).
- [36] Y. Qiao, C. Xiong, Z. Liu, and Z. Liu. 2019. Understanding the behaviors of bert in ranking. (2019). <https://arxiv.org/abs/1904.07531> arXiv: 1904.07531 [cs.LG].
- [37] N. Raifer, F. Raiber, M. Tennenholtz, and O. Kurland. 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 (SIGIR '17)*. Association for Computing Machinery, Shinjuku, Tokyo, Japan, 465–474. ISBN: 9781450350228. doi: [10.1145/3077136.3080785](https://doi.org/10.1145/3077136.3080785).
- [38] S. E. Robertson. 1977. The probability ranking principle in ir. *Journal of Documentation*, 33, 4, 294–304. doi: [10.1108/eb026647](https://doi.org/10.1108/eb026647).
- [39] F. Schroff, D. Kalenichenko, and J. Philbin. 2015. Facenet: a unified embedding for face recognition and clustering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. (June 2015), 815–823.
- [40] F. Yao, C. Li, D. Nekipelov, H. Wang, and H. Xu. 2023. How bad is top-K recommendation under competing content creators? In *Proceedings of the 40th International Conference on Machine Learning (Proceedings of Machine Learning Research)*. A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, and J. Scarlett, (Eds.) Vol. 202. PMLR, (July 2023), 39674–39701. <https://proceedings.mlr.press/v202/yao23b.html>.
- [41] F. Yao, C. Li, D. Nekipelov, H. Wang, and H. Xu. 2024. Human vs. generative AI in content creation competition: symbiosis or conflict? In *Proceedings of the 41st International Conference on Machine Learning (Proceedings of Machine Learning Research)*. Vol. 235. PMLR, (July 2024), 56885–56913. <https://proceedings.mlr.press/v235/yao24b.html>.
- [42] F. Yao, C. Li, K. A. Sankararaman, Y. Liao, Y. Zhu, Q. Wang, H. Wang, and H. Xu. 2023. Rethinking incentives in recommender systems: are monotone rewards always beneficial? In *Advances in Neural Information Processing Systems*. A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, (Eds.) Vol. 36. Curran Associates, Inc., 74582–74601. https://proceedings.neurips.cc/paper_files/paper/2023/file/ebfabf372037aaa4a8d92c9b457ece3e-Paper-Conference.pdf.
- [43] F. Yao et al. 2024. User welfare optimization in recommender systems with competing content creators. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '24)*. Association for Computing Machinery, Barcelona, Spain, 3874–3885. ISBN: 9798400704901. doi: [10.1145/3637528.3672021](https://doi.org/10.1145/3637528.3672021).
- [44] G. Yu, W. Tang, S. Narayanan, and C.-J. Ho. 2023. Encoding human behavior in information design through deep learning. In *Advances in Neural Information Processing Systems*. A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, (Eds.) Vol. 36. Curran Associates, Inc., 7506–7528. https://proceedings.neurips.cc/paper_files/paper/2023/file/17d0a21da4ec2c12b4f07fa2e34e4d6c-Paper-Conference.pdf.
- [45] J. Zhan, J. Mao, Y. Liu, M. Zhang, and S. Ma. 2020. Learning to retrieve: how to train a dense retrieval model effectively and efficiently. (2020). <https://arxiv.org/abs/2010.10469> arXiv: 2010.10469 [cs.LG].
- [46] W. X. Zhao, J. Liu, R. Ren, and J.-R. Wen. 2024. Dense text retrieval based on pretrained language models: a survey. *ACM Trans. Inf. Syst.*, 42, 4, Article 89, (Feb. 2024), 60 pages. doi: [10.1145/3637870](https://doi.org/10.1145/3637870).

A Equilibrium Analysis with the Euclidean Norm

In this section, we provide an equilibrium analysis of publishers' games that are based on the Euclidean norm, that is, publishers' games with $d(x, y) = \frac{1}{\sqrt{k}} \|x - y\|_2$. For each player $i \in N$, we define L_i as the line segment that connects the points x_0^i (her initial document) and x^* (the information need). Formally, L_i is defined as $L_i := \{tx^* + (1-t)x_0^i : t \in [0, 1]\}$. We use the notation $P_{L_i}(x_i)$ for the projection of x_i on L_i . We will make use of the following definitions:

DEFINITION 13. We say that a strategy $x_i \in X_i$ **strictly dominates** strategy $x'_i \in X_i$ (or: x'_i is **strictly dominated by** x_i) if for every strategy profile of the other publishers $x_{-i} \in X_{-i}$, x_i is a better response than x'_i with respect to x_{-i} . A strategy $x_i \in X_i$ is called a **strictly dominant strategy** if it dominates any other strategy of player i . A strategy $x_i \in X_i$ is called a **strictly dominated strategy** if it is strictly dominated by another strategy of player i .

Note that a strictly dominated strategy is never played in a pure Nash equilibrium. Interestingly, in many publishers' games, a significant number of strategies are strictly dominated.

LEMMA 5. Let d be a function of the form $d(a, b) = \frac{\|a-b\|^p}{C}$ for some $C > 0$ and $p \geq 1$, where $\|\cdot\|$ is the Euclidean (L_2) norm. Let G be a publishers' game with d as its semi-metric and whose ranking function r satisfies that $r_i(x)$ is

decreasing (not necessarily strictly) in $d^*(x_i)$. Then, for every player $i \in N$, any strategy $x_i \notin L_i$ is strictly dominated by $P_{L_i}(x_i)$.

PROOF. We use the fact that since L_i is convex and closed, $P_{L_i}(x_i)$ is well defined and unique, and the following holds [11]:

$$\forall y \in L_i, \langle x_i - P_{L_i}(x_i), y - P_{L_i}(x_i) \rangle \leq 0 \quad (39)$$

where $\langle \cdot, \cdot \rangle$ stands for the standard inner product.

Also, if $x_i \notin L_i$, then $P_{L_i}(x_i) \neq x_i$ so $\|x_i - P_{L_i}(x_i)\| > 0$. We get that for every $y \in L_i$:

$$\begin{aligned} \|x_i - y\|^2 &= \|x_i - P_{L_i}(x_i) + P_{L_i}(x_i) - y\|^2 \\ &= \|x_i - P_{L_i}(x_i)\|^2 + \|P_{L_i}(x_i) - y\|^2 + 2 \langle x_i - P_{L_i}(x_i), P_{L_i}(x_i) - y \rangle \\ &\stackrel{(39)}{\geq} \|x_i - P_{L_i}(x_i)\|^2 + \|P_{L_i}(x_i) - y\|^2 > \|P_{L_i}(x_i) - y\|^2 \end{aligned} \quad (40)$$

Hence in particular $\|x_i - x^*\| > \|P_{L_i}(x_i) - x^*\|$, and thus $d^*(x_i) > d^*(P_{L_i}(x_i))$, so by our assumption on r we have $r_i(x_i, x_{-i}) \leq r_i(P_{L_i}(x_i), x_{-i}) \forall x_{-i} \in X_{-i}$.

In the same way, $\|x_i - x_i^0\| > \|P_{L_i}(x_i) - x_i^0\|$, so $d_i^0(x_i) > d_i^0(P_{L_i}(x_i))$. By combining these two inequalities we get that for any $x_{-i} \in X_{-i}$:

$$\begin{aligned} u_i(x_i, x_{-i}) &= r_i(x_i, x_{-i}) - \lambda d_i^0(x_i) \\ &< r_i(P_{L_i}(x_i), x_{-i}) - \lambda d_i^0(P_{L_i}(x_i)) = u_i(P_{L_i}(x_i), x_{-i}) \end{aligned} \quad (41)$$

so by definition x_i is strictly dominated by $P_{L_i}(x_i)$ in G . \square

Note that the condition about the ranking function in Lemma 5 is satisfied by all ranking functions discussed in the paper. The Lemma implies, in particular, that players would never play outside their line segment L_i in a PNE. Lemma 6, unlike Lemma 5, does not hold for the squared Euclidean norm, and this is a main reason for the difference between the games induced by these two semi-metrics.

LEMMA 6. Let G be a publishers' game with $d(x, y) = \frac{1}{\sqrt{k}}\|x - y\|_2$. Then, for every player $i \in N$, $\forall x_i \in L_i$:

$$d^*(x_i) + d_i^0(x_i) = d(x_i^i, x^*) \quad (42)$$

PROOF. Let G be a publishers' game with $d(x, y) = \frac{1}{\sqrt{k}}\|x - y\|_2$, and let $i \in N$, $x_i \in L_i$. By definition, $\exists t \in [0, 1]$ s.t. $x_i = tx^* + (1 - t)x_0^i$.

$$\begin{aligned} d(x_0^i, x^*) &= \frac{\|x_0^i - x^*\|_2}{\sqrt{k}} = t \cdot \frac{\|x^* - x_0^i\|_2}{\sqrt{k}} + (1 - t) \cdot \frac{\|x_0^i - x^*\|_2}{\sqrt{k}} \\ &= \frac{\|tx^* - tx_0^i + x_0^i - x_0^i\|_2}{\sqrt{k}} + \frac{\|(1 - t)x_0^i - (1 - t)x^*\|_2}{\sqrt{k}} \\ &= \frac{\|tx^* + (1 - t)x_0^i - x_0^i\|_2}{\sqrt{k}} + \frac{\|tx^* + (1 - t)x_0^i - x^*\|_2}{\sqrt{k}} \\ &= \frac{\|x_i - x_0^i\|_2}{\sqrt{k}} + \frac{\|x_i - x^*\|_2}{\sqrt{k}} = d_i^0(x_i) + d^*(x_i) \end{aligned} \quad (43)$$

\square

Note that the term $d(x_0^i, x^*)$ depends only on the game parameters and does not depend on the strategy player i chooses to play. The above two Lemmas allow us to perform a thorough analysis of PNEs in games induced by the Euclidean norm and the linear ranking function.

LEMMA 7. Let G be a publishers' game with $r = \hat{r}$ and $d(x, y) = \frac{\|x-y\|_2}{\sqrt{k}}$.

- If $\lambda n < 1$, then x^{eq} defined by $x_i^{eq} = x^*$, $\forall i \in N$ is the only PNE in G , and $\forall i \in N$, x^* is a strictly dominant strategy for player i .
- If $\lambda n = 1$, then x^{eq} is a PNE in G if and only if $\forall i \in N$, $x_i^{eq} \in L_i$.
- If $\lambda n > 1$, then $x^{eq} = x_0$ is the only PNE in G , and $\forall i \in N$, x_0^i is a strictly dominant strategy for player i .

PROOF. Let G be a publishers' game with $r = \hat{r}$ and $d(x, y) = \frac{\|x-y\|_2}{\sqrt{k}}$. First, we notice that:

$$u_i(x) = r_i(x) - \lambda d_i^0(x_i) = \frac{v_i(x)}{n} + \frac{1}{n} - \lambda d_i^0(x_i) = \left(\sum_{j \in N \setminus \{i\}} \frac{d^*(x_j)}{n(n-1)} \right) - \frac{d^*(x_i)}{n} + \frac{1}{n} - \lambda d_i^0(x_i) \quad (44)$$

Now, for every player $i \in N$ and $\forall x_{-i} \in X_{-i}$:

$$\begin{aligned} \operatorname{argmax}_{x_i \in X_i} \{u_i(x_i, x_{-i})\} &\stackrel{(1)}{=} \operatorname{argmax}_{x_i \in L_i} \{u_i(x_i, x_{-i})\} \stackrel{(2)}{=} \operatorname{argmax}_{x_i \in L_i} \left\{ -\frac{d^*(x_i)}{n} - \lambda d_0(x_i) \right\} \\ &= \operatorname{argmin}_{x_i \in L_i} \left\{ d^*(x_i) + \lambda n d_0(x_i) \right\} \stackrel{(3)}{=} \operatorname{argmin}_{x_i \in L_i} \left\{ d(x_0^i, x^*) - d_i^0(x_i) + \lambda n d_0(x_i) \right\} \\ &= \operatorname{argmin}_{x_i \in L_i} \left\{ (\lambda n - 1) d_i^0(x_i) \right\} = \begin{cases} \{x^*\}, & \text{if } \lambda n < 1 \\ L_i, & \text{if } \lambda n = 1 \\ \{x_0^i\}, & \text{if } \lambda n > 1 \end{cases} \end{aligned} \quad (45)$$

When transition (1) is based on Lemma 5, transition (2) is substitution and removal of terms that do not depend on x_i and transition (3) is based on Lemma 6. This completes the proof, since a profile x is a PNE if and only if for every player $i \in N$ it holds that $x_i \in \operatorname{argmax}_{x_i' \in X_i} \{u_i(x_i', x_{-i})\}$. □

The resulting PNEs are degenerate in the sense that small changes in the value of λ or n can cause a complete change in both the PNEs and the publishers' dominant strategies.

B The Linear Ranking Function: Welfare Measures in Equilibrium

Let G be a publishers' game with $r = \hat{r}$ (that is, the linear ranking function with slope $a = \frac{1}{n}$) and $d(x, y) = \frac{1}{k} \|x - y\|_2^2$. According to Lemma 3, G has exactly one PNE which is $x_i^{eq} = \frac{x^* + \lambda n x_0^i}{1 + \lambda n}$. We aim to analyze the effect of the game parameters on $\mathcal{U}(x^{eq})$, the publishers' welfare in equilibrium, and on $\mathcal{V}(x^{eq})$, the users' welfare in equilibrium. We focus our analysis on the effect of λ , but a similar analysis can be made for other game parameters. Let us find a closed form for $\mathcal{U}(x^{eq})$ and for $\mathcal{V}(x^{eq})$. To this end, we denote by $d^i := d(x_0^i, x^*)$ the distance of publisher i 's initial document from the information need. We also denote by $d^{-i} := \frac{1}{n-1} \left(\sum_{j \in N \setminus \{i\}} d^j \right)$ the average of this distances for all publishers but i and by $\bar{d} := \frac{1}{n} \sum_{i=1}^n d^i$ the average of those distances for all publishers. Lastly, let us denote $d' := \frac{1}{n} \sum_{i=1}^n d^i \cdot (d^{-i} - d^i)$. Let us now perform some simple calculations and substitutions. Firstly, the distances of player i 's equilibrium strategy from the information need and from her initial document

are given by:

$$\begin{aligned} d^*(x_i^{eq}) &= \frac{1}{k} \|x_i^{eq} - x^*\|_2^2 = \frac{1}{k} \left\| \frac{x^* + \lambda n x_0^i}{1 + \lambda n} - x^* \right\|_2^2 \\ &= \frac{(\lambda n)^2}{k(1 + \lambda n)^2} \|x_0^i - x^*\|_2^2 = \frac{(\lambda n)^2}{(1 + \lambda n)^2} d^i \end{aligned} \quad (46)$$

$$\begin{aligned} d_i^0(x_i^{eq}) &= \frac{1}{k} \|x_i^{eq} - x_0^i\|_2^2 = \frac{1}{k} \left\| \frac{x^* + \lambda n x_0^i}{1 + \lambda n} - x_0^i \right\|_2^2 \\ &= \frac{1}{k(1 + \lambda n)^2} \|x^* - x_0^i\|_2^2 = \frac{1}{(1 + \lambda n)^2} d^i \end{aligned} \quad (47)$$

Player i 's linear relative relevance and player i 's probability to be ranked first are given by the following terms, respectively:

$$\begin{aligned} v_i(x^{eq}) &= \frac{1}{n-1} \left(\sum_{j \in N \setminus \{i\}} d^*(x_j^{eq}) \right) - d^*(x_i^{eq}) \\ &= \frac{1}{n-1} \left(\sum_{j \in N \setminus \{i\}} \frac{(\lambda n)^2 d^j}{(1 + \lambda n)^2} \right) - \frac{(\lambda n)^2 d^i}{(1 + \lambda n)^2} \\ &= \frac{(\lambda n)^2}{(1 + \lambda n)^2} (d^{-i} - d^i) \end{aligned} \quad (48)$$

$$\hat{r}_i(x^{eq}) = \frac{1}{n} v_i(x^{eq}) + \frac{1}{n} = \frac{\lambda^2 n}{(1 + \lambda n)^2} (d^{-i} - d^i) + \frac{1}{n} \quad (49)$$

Now we can substitute what we found to find the equilibrium utility of player i :

$$\begin{aligned} u_i(x^{eq}) &= \hat{r}_i(x^{eq}) - \lambda d_i^0(x_i^{eq}) \\ &= \frac{\lambda^2 n}{(1 + \lambda n)^2} (d^{-i} - d^i) + \frac{1}{n} - \frac{\lambda d^i}{(1 + \lambda n)^2} \\ &= \frac{\lambda^2 n (d^{-i} - d^i) - \lambda d^i}{(1 + \lambda n)^2} + \frac{1}{n} \end{aligned} \quad (50)$$

We are now ready to substitute the previous equations in order to find a closed form for the welfare measures in equilibrium. The publishers' welfare in equilibrium is simply:

$$\mathcal{U}(x^{eq}) = 1 - \lambda \sum_{i=1}^n d_i^0(x_i^{eq}) = 1 - \lambda \sum_{i=1}^n \frac{1}{(1 + \lambda n)^2} d^i = 1 - \frac{\lambda n}{(1 + \lambda n)^2} \bar{d} \quad (51)$$

and the users' welfare is given by:

$$\begin{aligned} \mathcal{V}(x^{eq}) &= 1 - \sum_{i=1}^n \hat{r}_i(x^{eq}) \cdot d^*(x_i^{eq}) = 1 - \sum_{i=1}^n \left(\frac{\lambda^2 n}{(1 + \lambda n)^2} (d^{-i} - d^i) + \frac{1}{n} \right) \cdot \frac{(\lambda n)^2}{(1 + \lambda n)^2} d^i \\ &= 1 - \frac{(\lambda n)^2}{(1 + \lambda n)^2} \cdot \frac{1}{n} \sum_{i=1}^n d^i - \frac{(\lambda n)^4}{(1 + \lambda n)^4} \cdot \frac{1}{n} \sum_{i=1}^n d^i \cdot (d^{-i} - d^i) \\ &= 1 - \frac{(\lambda n)^2}{(1 + \lambda n)^2} \cdot \bar{d} - \frac{(\lambda n)^4}{(1 + \lambda n)^4} \cdot d' \end{aligned} \quad (52)$$

We will now examine the impact of λ on both publishers' and users' welfare in equilibrium. This parameter represents the publishers' incentive to adhere to their initial document. Therefore, understanding the impact of λ is vital for system designers striving to optimize welfare measures.

B.1 Publishers' Welfare Analysis

We start by finding the partial derivative of $\mathcal{U}(x^{eq})$ with respect to λ .

$$\frac{\partial \mathcal{U}(x^{eq})}{\partial \lambda} = -\frac{n(1+\lambda n)^2 - 2\lambda n^2(1+\lambda n)}{(1+\lambda n)^4} \cdot \bar{d} = \frac{(1+\lambda n) - 2\lambda n}{(1+\lambda n)^3} \cdot \bar{d}n = \frac{1-\lambda n}{(1+\lambda n)^3} \cdot \bar{d}n \quad (53)$$

The denominator is always positive, therefore $\frac{\partial \mathcal{U}(x^{eq})}{\partial \lambda}$ is positive when $\lambda < \frac{1}{n}$, zero when $\lambda = \frac{1}{n}$ and negative when $\lambda > \frac{1}{n}$. We get that $\lambda = \frac{1}{n}$ is a global minimum point of $\mathcal{U}(x^{eq})$.

B.2 Users' Welfare Analysis

We start by finding the partial derivative of $\mathcal{V}(x^{eq})$ with respect to λ . Recall that the users' welfare in equilibrium is given by:

$$\mathcal{V}(x^{eq}) = 1 - \frac{(\lambda n)^2}{(1+\lambda n)^2} \cdot \bar{d} - \frac{(\lambda n)^4}{(1+\lambda n)^4} \cdot d' \quad (54)$$

For ease of derivation, let us denote $f(t) := 1 - t\bar{d} - t^2d'$ and $g(\lambda) := \frac{(\lambda n)^2}{(1+\lambda n)^2}$. Let us write the derivatives of these two functions: $f'(t) = -\bar{d} - 2td'$, $g'(\lambda) = \frac{2\lambda n^2}{(1+\lambda n)^3}$. We can now write:

$$\mathcal{V}(x^{eq}) = f(g(\lambda)) \quad (55)$$

and thus by the chain rule:

$$\begin{aligned} \frac{\partial \mathcal{V}(x^{eq})}{\partial \lambda} &= f'(g(\lambda))g'(\lambda) = (-\bar{d} - 2\frac{(\lambda n)^2}{(1+\lambda n)^2}d') \cdot \frac{2\lambda n^2}{(1+\lambda n)^3} \\ &= \frac{2\lambda n^2}{(1+\lambda n)^5} \left(-\bar{d}(1+2\lambda n + \lambda^2 n^2) - 2\lambda^2 n^2 d' \right) \\ &= \underbrace{\frac{2\lambda n^2}{(1+\lambda n)^5}}_{:= C(\lambda)} \underbrace{\left(-n^2(\bar{d} + 2d')\lambda^2 - 2n\bar{d}\lambda - \bar{d} \right)}_{:= h(\lambda)} \end{aligned} \quad (56)$$

We now aim to determine the range of λ values for which the derivative is positive and the range for which it is negative. As $C(\lambda) > 0$ for all $\lambda > 0$ (recall we are only interested in positive λ values), the sign of the derivative is simply the sign of $h(\lambda)$. We define the condition $\bar{d} + 2d' \geq 0$ as the regularity condition and divide our analysis into two cases:

- **Case I - the regularity condition holds** ($\bar{d} + 2d' \geq 0$): In this case, $h(\lambda) \leq 0 \forall \lambda > 0$. So $\frac{\partial \mathcal{V}(x^{eq})}{\partial \lambda} \leq 0 \forall \lambda > 0$ and therefore in this case the users' welfare in equilibrium is monotonically decreasing in λ .
- **Case II - the regularity condition does not hold** ($\bar{d} + 2d' < 0$): In this case $h(\lambda)$ is a strictly convex quadratic function, so it suffices for us to find its roots in order to know its positivity and negativity ranges. By the quadratic formula, h 's roots are:

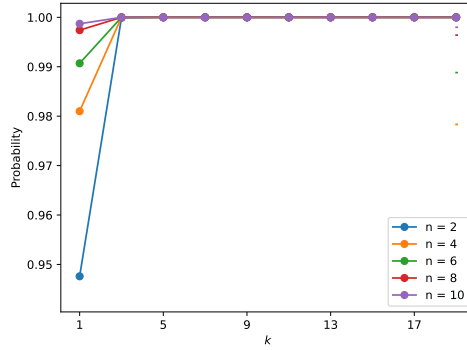


Fig. 5. Probability of regularity condition with independently and uniformly distributed x_0 and x^* .

$$\begin{aligned} \lambda_{1,2} &= \frac{2n\bar{d} \pm \sqrt{(-2n\bar{d})^2 - 4(-n^2(\bar{d} + 2d'))(-\bar{d})}}{-2n^2(\bar{d} + 2d')} \\ &= \frac{2n\bar{d} \pm \sqrt{-8n^2\bar{d}d'}}{-2n^2(\bar{d} + 2d')} = \frac{\bar{d} \pm \sqrt{-2\bar{d}d'}}{-n(\bar{d} + 2d')} \end{aligned} \quad (57)$$

Let us show that the term inside the square root is strictly positive. First, $\bar{d} \geq 0$ as the average of non-negative terms. $\bar{d} = 0$ would imply that $d^i = 0 \forall i$ and then we get $d' = 0$ which contradicts $\bar{d} + 2d' < 0$. Therefore, $\bar{d} > 0$. Additionally, from $\bar{d} + 2d' < 0$ we get $\bar{d} < -2d'$ and thus also $0 < -2d'$. So $-2\bar{d}d' > 0$. Now we know that both λ_1 and λ_2 are well-defined and that h has two distinct roots. Furthermore, we notice that

$$\lambda_1 := \frac{\bar{d} + \sqrt{-2\bar{d}d'}}{-n(\bar{d} + 2d')} > 0 \quad (58)$$

since the numerator is strictly positive ($\bar{d} > 0$) and the denominator is strictly positive ($\bar{d} + 2d' < 0$). As for the other root,

$$\lambda_2 := \frac{\bar{d} - \sqrt{-2\bar{d}d'}}{-n(\bar{d} + 2d')} < 0 \quad (59)$$

because $\lambda_2 < 0 \iff \bar{d} - \sqrt{-2\bar{d}d'} < 0 \iff \bar{d} < \sqrt{-2\bar{d}d'} \iff_{(1)} \bar{d}^2 < -2\bar{d}d' \iff \bar{d}^2 + 2\bar{d}d' < 0 \iff_{(2)} \bar{d} + 2d' < 0$ where both (1) and (2) can be justified by $\bar{d} > 0$, which we already showed.

Summing up this case, we got that $h(\lambda)$ is a strictly convex quadratic function with roots $\lambda_2 < 0 < \lambda_1$. So $h(\lambda)$ is negative when $0 < \lambda < \lambda_1$ and positive when $\lambda_1 < \lambda$. Therefore, these are also the negativity and positivity ranges of $\frac{\partial \mathcal{V}(x^{eq})}{\partial \lambda}$. Thus, λ_1 is a global minimum point of $\mathcal{V}(x^{eq})$.

REMARK 1. The regularity condition $\bar{d} + 2d' \geq 0$ is equivalent to requiring that the unbiased sample variance of $\{d^i\}_{i \in N}$ is no greater than half the empirical mean \bar{d} . This follows from the fact that $d' = \frac{n}{n-1} \left(\bar{d}^2 - \frac{1}{n} \sum_{i=1}^n (d^i)^2 \right)$, which shows that d' is equal to minus the unbiased sample variance of $\{d^i\}_{i \in N}$.

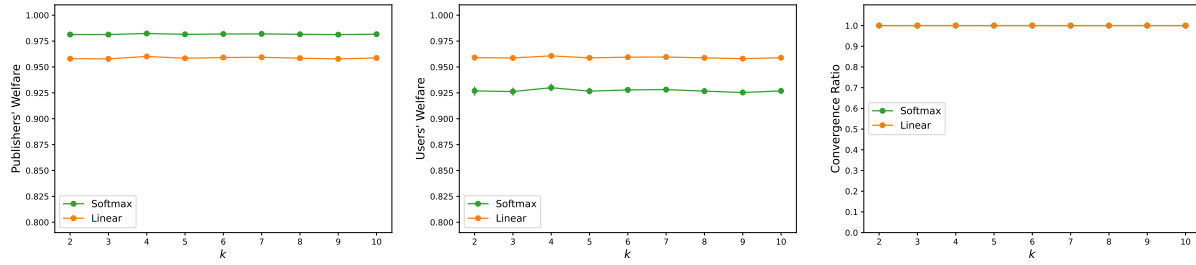


Fig. 6. The effect of k in the softmax ($\beta = 1$) and the linear ($a = \frac{1}{n} = 0.5$) ranking functions with $\lambda = 0.5$ and $n = 2$.

In Figure 5 we estimate the probability for the regularity condition for different n and k values with independently and uniformly distributed x_0 and x^* .

To sum up, in this section we have demonstrated how the closed form of the equilibrium allows us to easily analyze the effect of λ on the welfare measures in the PNE. A similar analysis can be performed on the effect of the other game parameters.

C The Effect of the Embedding Space Dimension

In this section, we discuss the effect of the embedding space dimension, k . This examination is more complicated than examining the effect of λ since the time complexity of the simulation increases exponentially with k . We first discuss the effect of k when using the softmax and the linear ranking functions we explored in §6 and then present two approaches to assess this effect when using the PRP ranking function, a direct approach and a dynamics-based approach. Finally, we discuss how the results of this section affect the comparison between the different ranking functions. Note that $k = 1$ is essentially different than other k values since it is the only case in which the possible movement directions in the Discrete Better Response Dynamics Simulation include all possible directions in the embedding space. This attribute, together with the fact that $k = 1$ is not a realistic case regardless, leads us to omit it from our analysis.

C.1 Direct Examination of the Effect of k on the Softmax and the Linear Ranking Functions

We performed 500 simulations for each k value to directly test the effect of k on the welfare measures and on the convergence ratio under the softmax and the linear ranking functions.²³ The results of these simulations are presented in Figure 6.

We can see that the convergence ratios of both ranking functions remain 1 for all tested k values. Furthermore, Figure 6 evidences that when using the softmax or the linear ranking functions k has a negligible effect on both the publishers' and the users' welfare. We observed similar results for other λ values and different ranking functions hyperparameters. Therefore, it can be said that the results in §6 regarding those functions can be extrapolated to any low-mid k value. Examining high values of k is highly impractical using this kind of simulation due to the exponential growth of the simulation running time as k increases. However, note that all our theoretical results regarding the linear ranking function hold for any k . In particular, the convergence of any ε -better response dynamics to a ε -PNE is guaranteed, and thus the convergence ratio of the linear ranking function is guaranteed to be 1 for any k .

²³In these simulations we used the same simulations' parameters as in the simulations presented in §6.

$S = \{2^{-6}, 2^{-5}, \dots, 2^{-2}, 0.5, 0.6, \dots, 1\}$, $\mathcal{D} = \{\frac{d}{\|d\|_2} : d \in \{-1, 0, 1\}^k \setminus \{\vec{0}\}\}$, $T = 1000$, $M = 900$ and $\varepsilon = 10^{-6}$, with a bootstrap sample size of $B = 500$.

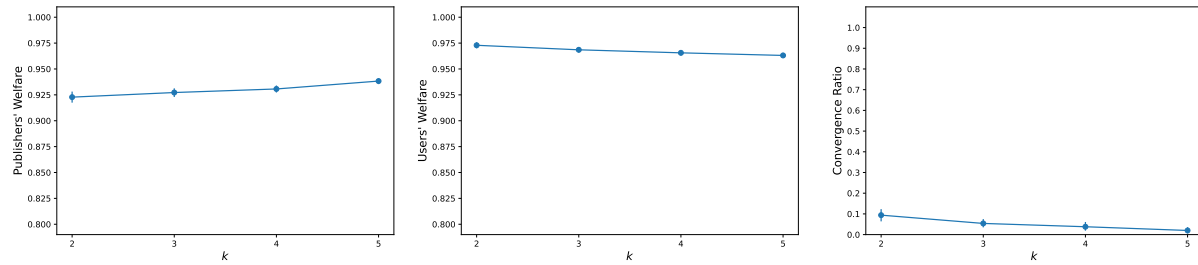


Fig. 7. The effect of k in the PRP ranking function with $\lambda = 0.5$ and $n = 2$.

Furthermore, in Appendix A we presented Lemma 5, which determines that in games induced by various ranking functions, including the softmax and the linear ranking functions, publishers should never play a strategy (document) that is not on the line segment that connects x^* with their initial document. This result shows us that the nature of the dynamics in publishers' games induced by these ranking functions is highly one-dimensional. That is, no matter how large k is, each player still plays roughly on her one-dimensional line segment. Thus, we conjecture that the softmax ranking function will have a convergence ratio of 1 for all k . Furthermore, based on this argument, we can also conjecture that k will have a negligible effect on the welfare measures of both the linear and softmax ranking functions.

C.2 Direct Examination of the Effect of k on the PRP Ranking Function

We now turn to discuss the effect of k in PRP-induced publishers' games, which introduces an increase in analytical complexity. In the Discrete Better Response Dynamics Simulation, the number of movement directions, and hence the simulation running time, grows exponentially with k . This challenge is further amplified when simulating the PRP ranking function, as the simulations frequently fail to converge. This makes simulations of the PRP ranking function with high values of k highly impractical. Figure 7 presents the results for a few low values of k .²⁴

Figure 7 may indicate that, as is the case for the softmax and the linear ranking functions, k has a minor effect on the publishers' and users' welfare when using the PRP ranking function. We can see that as k increases, the publishers' welfare slightly increases and the users' welfare slightly decreases. As mentioned, due to the structure of the simulation, we can't examine the welfare measures with higher k values, but we conjecture that the observed trends will continue until a certain limit.

Regarding the convergence ratio, Figure 7 might suggest that it generally decreases with k and that the decrease becomes more moderate as k increases. As we only tested low values of k , it remains uncertain whether this trend holds for higher values of k as well. Therefore, we employ an alternative approach to approximate the effect of k on the convergence ratio of the PRP ranking function.

C.3 Dynamics-Based Examination of the Effect of k on the Convergence Ratio of the PRP Ranking Function

In §6.4 we observed the two primary types of dynamics induced by the PRP ranking function, the uncommon type where the simulation converges already in the first timestep, and the common type where the publishers

²⁴In this examination we ran 500 simulations for each k value, with $S = \{2^{-6}, 2^{-5}, \dots, 2^{-2}, 0.5, 0.6, \dots, 1\}$, $\mathcal{D} = \{\frac{d}{\|d\|_2} : d \in \{-1, 0, 1\}^k \setminus \{\vec{0}\}\}$, $T = 5000$, $M = 4500$ and $\varepsilon = 10^{-6}$, with a bootstrap sample size of $B = 500$. We chose high T and M values since the cycle length of the pseudo-periodic dynamics of the PRP ranking function (which we discussed in §6.4) increases as k increases.

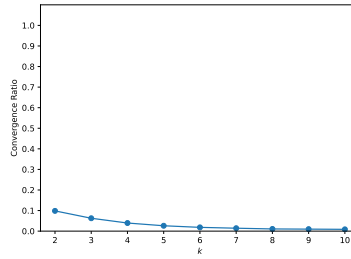


Fig. 8. The effect of k on the PRP ranking function's convergence ratio with $\lambda = 0.5$ and $n = 2$.

reach a pseudo-periodic behavior. Based on these results, we can use a new method to approximate the effect of k on the convergence ratio of the PRP ranking function. Instead of running a whole simulation each time, we can determine convergence at the beginning of the simulation. If the dynamics did not converge in the first timestep, we can deduce that the simulation will not converge. Figure 8 presents the convergence ratio measured using this estimation method.²⁵

Figure 8 shows that the convergence ratio of the PRP ranking function generally decreases as k increases and that the decrease becomes more moderate as k increases. In §6.4, we saw that the case when PRP dynamics converge is when the publisher whose initial document is farther from x^* cannot bypass the other publisher in a single step. Therefore, a possible explanation for the decrease in the convergence ratio is that a larger k decreases the likelihood of this case. Although Figure 8 provides results only for $\lambda = 0.5$, similar trends were observed for other values of λ . We conjecture that these trends continue at higher k values.

To conclude this section, let us revisit the trends we discovered in light of the results we have already seen in §6. From this perspective, we can see how an increase in k emphasizes the consequences of the instability of the PRP ranking function. The already low convergence ratio at $k = 2$ further diminishes with increasing k . Simultaneously, the users' welfare, a measure that the PRP is designed to maximize in the short term, also experiences a minor decline as k increases. In contrast, the examination of different k values underscores the stability of both the softmax and the linear ranking functions, maintaining a convergence ratio of 1 and roughly identical welfare measures across all tested k values.

D Additional Experiments

In this section, we provide and discuss additional experimental results. These experiments aim to provide some preliminary insights regarding settings beyond the ones presented in the body of the paper. In all simulations in this section, we used the same simulation parameters as in the simulations presented in §6.

D.1 Beyond the Two-Publisher Case

Figure 9 depicts the effect of n on our three evaluation criteria (publishers' welfare, users' welfare, and convergence ratio), under the three ranking schemes. Importantly, the slope of the linear ranking function decreases as n grows large (that is, $a = \frac{1}{n}$). As a result, the performance of the linear ranking function deteriorates in terms of users' welfare as the number of publishers grows. Note that for lower values of n , the softmax ranking function

²⁵In this special method we have conducted 10,000 simulations with $S = \{2^{-6}, 2^{-5}, \dots, 2^{-2}, 0.5, 0.6, \dots, 1\}$, $\mathcal{D} = \left\{ \frac{d}{\|d\|_2} : d \in \{-1, 0, 1\}^k \setminus \{\vec{0}\} \right\}$ and $\varepsilon = 10^{-6}$, and used a bootstrap sample size of $B = 10,000$. We have conducted a large number of simulations, since the measurement of the convergence ratio involves a lot of noise, and we could not use the same group of randomized x_0 and x^* for all the k values due to the different vector dimensions. Conducting a large number of simulations was made possible thanks to our method of running each simulation for a single timestep.

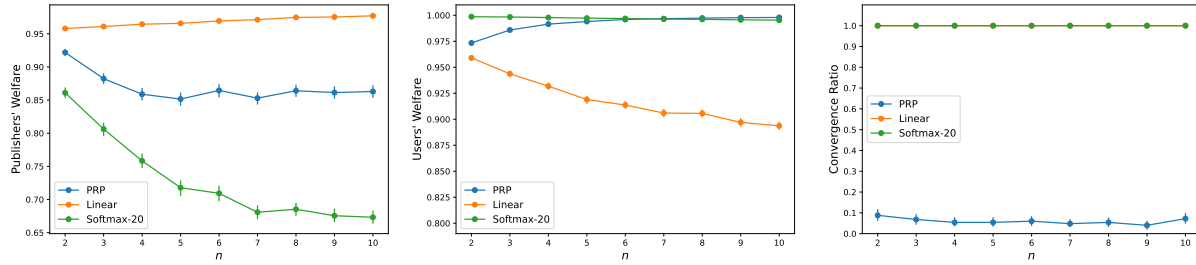


Fig. 9. Comparison of the PRP, softmax ($\beta = 20$) and linear ($a = \frac{1}{n}$) ranking functions across different n values with constant $\lambda = \frac{1}{2}$. The linear and the softmax ranking functions had a constant convergence ratio of 1.

(with inverse temperature $\beta = 20$) outperforms the PRP ranking in terms of users' welfare. However, as n grows, the users' welfare under PRP ranking improves until it reaches the users' welfare under softmax ranking. That being said, it might be the case that adaptively choosing β as a function of n can yield a consistent improvement over the PRP in terms of users' welfare. The stability of the linear and softmax schemes (as well as the instability of the PRP) is also preserved.

The decrease in users' welfare under linear ranking is expected, since the slope $a = \frac{1}{n}$ decreases as n increases, limiting the extent to which the ranker can incentivize publishers towards relevant content (see the discussion in §5.3). Having said that, we highlight that the case of small n (even $n = 2$) is widely observed when publishers compete in a niche domain. This also aligns with the assumption of a single, known x^* – publishers specialize in a specific domain and know precisely the users' preferences and required content. For higher values of n , the softmax ranking function (with an appropriate β) provides a better alternative in terms of users' welfare since β can be taken arbitrarily large.

D.2 Beyond the Uniform i.i.d. Instance Distribution Case

We now consider a different family of instance distributions, on which we evaluate the three ranking schemes. These distributions diverge from the assumptions of independence between the initial documents and the information need, as well as the assumption of uniform marginal distributions. Specifically, we sample from a multivariate normal distribution, conditioned to lie within $[0, 1]^k$. For each coordinate $j \in [k] = [2]$, the j -th entries of the $n = 2$ publishers' initial documents and the information need are sampled from the conditional multivariate normal distribution with mean $(0.5, 0.5, 0.5)$, and covariance matrix of:

$$\begin{array}{c|ccc} & [x_0^1]_j & [x_0^2]_j & [x^*]_j \\ \hline [x_0^1]_j & 1 & \rho_1 & \rho_2 \\ [x_0^2]_j & \rho_1 & 1 & \rho_1\rho_2 \\ [x^*]_j & \rho_2 & \rho_1\rho_2 & 1 \end{array}$$

Note that while entries in different coordinates remain independent, there is a correlation between the initial documents and the information need within each coordinate. When $\rho_1 = \rho_2 = 0$, this distribution corresponds to a conditional multivariate normal distribution with the initial documents and the information need being pairwise independent. In Figure 10a we present the trends of publishers' welfare, users' welfare, and convergence ratio as a function of λ for $\rho_1 = \rho_2 = 0$. Note that the trends are as in the uniform case presented in §6.2.

Next, we introduce three additional cases, in which there is a correlation between the initial documents and the information need. The first case, presented in Figure 10b, corresponds to $\rho_1 = \frac{1}{2}, \rho_2 = 0$. Intuitively, this means that the publishers' initial documents are positively correlated. The second case, presented in Figure 10c,

corresponds to $\rho_1 = 0, \rho_2 = \frac{1}{2}$, which means that the initial document of publisher 1 will be generally closer to the information need. The third case, presented in Figure 10d, corresponds to $\rho_1 = \frac{1}{2}, \rho_2 = \frac{1}{2}$, which means that both trends of the first two cases occur, and as a result there is some weaker correlation between the initial document of publisher 2 and the information need.

Notice that in all three cases, the trends again agree with the ones presented in the paper. While these results are preliminary, we believe similar trends are preserved in many other natural instance distributions as well.

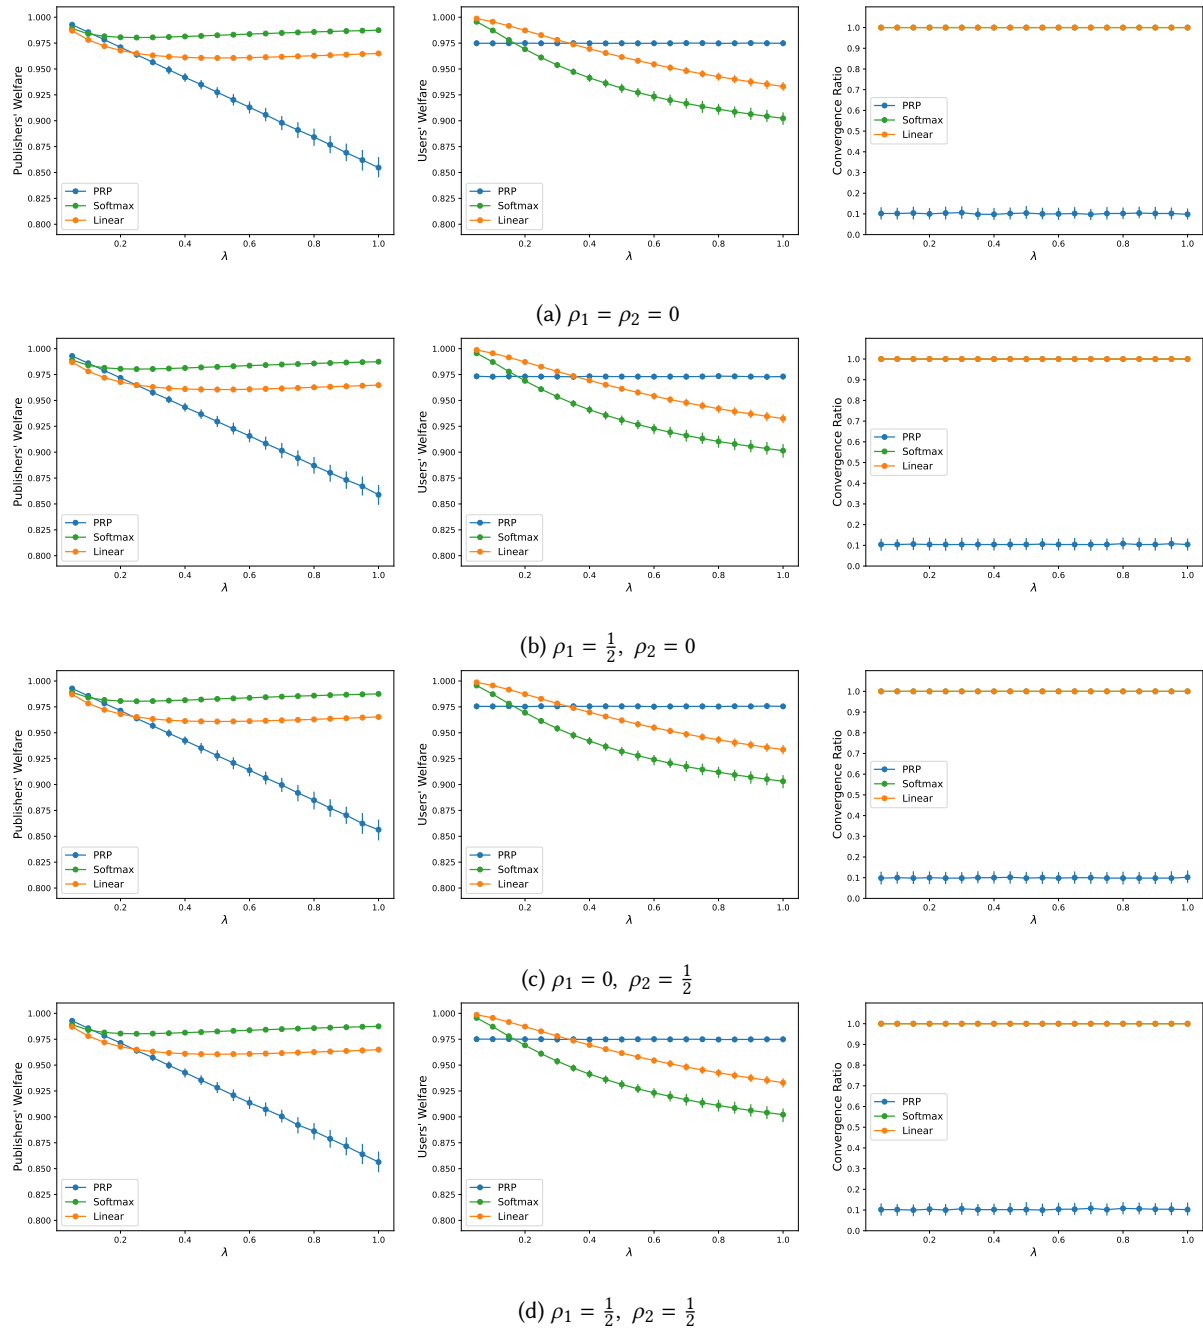


Fig. 10. Comparison of the PRP, softmax ($\beta = 1$) and linear ($a = \frac{1}{n}$) ranking functions with initial documents and the information need being distributed according to the conditional multivariate normal distribution. The parameters of the distribution are provided below each figure. In all cases, The linear and the softmax ranking functions had a constant convergence ratio of 1.

E Omitted Proofs of the Auxiliary Lemmas

PROOF OF LEMMA 1. Assume by contradiction that there exists $i \in N$ such that $x_i^{eq} \notin \{x^*, x_0^i\}$. By Lemma 5 we get that $x_i^{eq} \in L_i := \{tx^* + (1-t)x_0^i : t \in [0, 1]\}$. Thus $x_0^i \neq x^*$.

Recall the notation $\mu(x) := \operatorname{argmin}_{i \in N} d^*(x_i)$ for $x \in X$. Let us divide the proof into the following cases:

- (1) If $i \notin \mu(x^{eq})$ then $u_i(x^{eq}) = -\lambda d_i^0(x_i^{eq})$. Since $x_i^{eq} \neq x_0^i$ we get that $d_i^0(x_i^{eq}) > 0$ and hence $u_i(x^{eq}) < 0$. On the other hand, $u_i(x_0^i, x_{-i}^{eq}) = r_i(x_0^i, x_{-i}^{eq}) - 0 \geq 0$. Hence $u_i(x_0^i, x_{-i}^{eq}) > u_i(x^{eq})$ which is a contradiction to the fact that x^{eq} is a PNE.
- (2) If $\mu(x^{eq}) = \{i\}$ then $d^*(x_i^{eq}) < d^*(x_j^{eq}) \forall j \neq i$. Let us denote $\varepsilon = \frac{1}{2}(\min_{j \neq i} d^*(x_j^{eq}) - d^*(x_i^{eq})) > 0$. Since d is continuous, d^* is continuous and hence there exists a $\delta > 0$ such that if $\|x_i - x_i^{eq}\|_2 < \delta$ then $d^*(x_i) < d^*(x_i^{eq}) + \varepsilon$. Therefore, it holds in particular for

$$\hat{x}_i := x_i^{eq} + \frac{\delta}{2\|x_i^{eq} - x_0^i\|_2}(x_0^i - x_i^{eq}) \quad (60)$$

that $d^*(\hat{x}_i) < d^*(x_i^{eq}) + \varepsilon$ and hence $d^*(\hat{x}_i) < \min_{j \neq i} d^*(x_j^{eq})$. Therefore $\mu(\hat{x}_i, x_{-i}^{eq}) = \{i\}$ and hence $r_i^*(\hat{x}_i, x_{-i}^{eq}) = 1$. We get that:

$$\begin{aligned} u_i(\hat{x}_i, x_{-i}^{eq}) &= 1 - \lambda d_i^0(\hat{x}_i) = 1 - \lambda \frac{1}{k} \|\hat{x}_i - x_0^i\|_2^2 = 1 - \lambda \frac{1}{k} \left\| (x_i^{eq} - x_0^i) \left(1 - \frac{\delta}{2\|x_i^{eq} - x_0^i\|_2}\right) \right\|_2^2 \\ &= 1 - \lambda \frac{1}{k} \left(1 - \frac{\delta}{2\|x_i^{eq} - x_0^i\|_2}\right)^2 \|x_i^{eq} - x_0^i\|_2^2 > 1 - \lambda \frac{1}{k} \|x_i^{eq} - x_0^i\|_2^2 \\ &= 1 - \lambda d_i^0(x_i^{eq}) = u_i(x^{eq}) \end{aligned} \quad (61)$$

which is in contradiction to the fact that x^{eq} is a PNE.

- (3) If $i \in \mu(x^{eq})$ and $|\mu(x^{eq})| \geq 2$:

d is continuous and hence d_i^0 is continuous. Therefore there exists a $\delta > 0$ such that if $\|x_i - x_i^{eq}\|_2 < \delta$ then $d_i^0(x_i) < d_i^0(x_i^{eq}) + \frac{1}{2\lambda}$. Therefore, it holds in particular for

$$\hat{x}_i := x_i^{eq} + \frac{\delta}{2\|x_i^{eq} - x^*\|_2}(x^* - x_i^{eq}) \quad (62)$$

that $d_i^0(\hat{x}_i) < d_i^0(x_i^{eq}) + \frac{1}{2\lambda}$.

Now, $d^*(\hat{x}_i) = \frac{1}{k} \|\hat{x}_i - x^*\|_2^2 = \frac{1}{k} \left\| (x_i^{eq} - x^*) \left(1 - \frac{\delta}{2\|x_i^{eq} - x^*\|_2}\right) \right\|_2^2 < \frac{1}{k} \|x_i^{eq} - x^*\|_2^2 = d^*(x_i^{eq})$. Therefore $\mu(\hat{x}_i, x_{-i}^{eq}) = \{i\}$ and hence $r_i(\hat{x}_i, x_{-i}^{eq}) = 1$. We get that:

$$\begin{aligned} u_i(\hat{x}_i, x_{-i}^{eq}) &= 1 - \lambda d_i^0(\hat{x}_i) > 1 - \lambda \left(d_i^0(x_i^{eq}) + \frac{1}{2\lambda} \right) \\ &= \frac{1}{2} - \lambda d_i^0(x_i^{eq}) \geq \frac{1}{|\mu(x^{eq})|} - \lambda d_i^0(x_i^{eq}) = u_i(x^{eq}) \end{aligned} \quad (63)$$

which is once again a contradiction to the fact that x^{eq} is a PNE. \square

PROOF OF LEMMA 2. If g is constant on $[0, 1]$ then $f_i \equiv \frac{1}{n} \forall i \in [n]$ and then f satisfies (12) trivially, since both sides of the equation are 0.

In the other direction, assume that f satisfies (12). Let us calculate the derivatives. It holds $\forall i, j \in [n]$ and $\forall x \in X$ that:

$$\frac{\partial f_i(x)}{\partial x_j} = \frac{-g(x_i)g'(x_j)}{\left(\sum_{k=1}^n g(x_k)\right)^2} \quad (64)$$

and

$$\frac{\partial^2 f_i(x)}{\partial x_i \partial x_j} = \frac{g'(x_i)g'(x_j) \left(2g(x_i) - \sum_{k=1}^n g(x_k)\right)}{\left(\sum_{k=1}^n g(x_k)\right)^3} \quad (65)$$

In the same way, we get:

$$\frac{\partial^2 f_j(x)}{\partial x_i \partial x_j} = \frac{g'(x_i)g'(x_j) \left(2g(x_j) - \sum_{k=1}^n g(x_k)\right)}{\left(\sum_{k=1}^n g(x_k)\right)^3} \quad (66)$$

So by (12) we get:

$$0 = \frac{\partial^2 f_i(x)}{\partial x_i \partial x_j} - \frac{\partial^2 f_j(x)}{\partial x_i \partial x_j} = \frac{g'(x_i)g'(x_j) (2g(x_i) - 2g(x_j))}{\left(\sum_{k=1}^n g(x_k)\right)^3} \quad (67)$$

Hence, $\forall i, j \in [n]$ and $\forall x \in [0, 1]^n$ it holds that:

$$g'(x_i) = 0 \vee g'(x_j) = 0 \vee g(x_i) = g(x_j) \quad (68)$$

Equivalently, (since $n \geq 2$):

$$g'(a) = 0 \vee g'(b) = 0 \vee g(a) = g(b) \quad \forall a, b \in [0, 1] \quad (69)$$

We proceed to show that (69) implies that g is constant on $[0, 1]$. We divide the proof into two cases:

- If there is no $a \in [0, 1]$ such that $g'(a) \neq 0$, then $g' \equiv 0$ on $[0, 1]$ and therefore g is constant on $[0, 1]$.
- Otherwise, let $a \in [0, 1]$ be such that $g'(a) \neq 0$. Without loss of generality, assume $g'(a) > 0$. By continuity of g' , there is a $\delta > 0$ such that $g'(x) > 0 \forall x \in U := (a - \delta, a + \delta) \cap [0, 1]$. So g is strictly increasing on U . Take $b \in U$. Since g is strictly increasing on U , $g(a) \neq g(b)$. Also, $g'(b) \neq 0$, $g'(a) \neq 0$, since they are both positive. This is a contradiction to (69). □

PROOF OF LEMMA 4. We have that:

$$\psi_i(d) = \frac{1}{n-1} \sum_{j \neq i} d_j - d_i = \frac{1}{n-1} \sum_{j=1}^n d_j - \frac{n}{n-1} d_i \quad (70)$$

Therefore:

$$\psi_i(d) = \psi_j(d) \iff \frac{n}{n-1} d_i = \frac{n}{n-1} d_j \iff d_i = d_j \quad (71)$$

□

Received 31 December 2024; revised 15 June 2025; accepted 24 June 2025