The Probability Ranking Principle is Not Optimal in Adversarial Retrieval Settings

Ran Ben Basat
Computer Science Department
Technion
Haifa, Israel
sran@cs.technion.ac.il

Moshe Tennenholtz
Faculty of Industrial Engineering and Management
Technion
Haifa, Israel
moshet@ie.technion.ac.il

Oren Kurland
Faculty of Industrial Engineering and Management
Technion
Haifa, Israel
kurland@ie.technion.ac.il

ABSTRACT

The probability ranking principle (PRP) — ranking documents in response to a query by their relevance probabilities — is the theoretical foundation of most ad hoc document retrieval methods. A key observation that motivates our work is that the PRP does not account for potential post-ranking effects, specifically, changes to documents that result from a given ranking. Yet, in adversarial retrieval settings such as the Web, authors may consistently try to promote their documents in rankings by changing them. We prove that, indeed, the PRP can be sub-optimal in adversarial retrieval settings. We do so by presenting a novel game theoretic analysis of the adversarial setting. The analysis is performed for different types of documents (single topic and multi topic) and is based on different assumptions about the writing qualities of documents' authors. We show that in some cases, introducing randomization into the document ranking function yields overall user utility that transcends that of applying the PRP.

Categories and Subject Descriptors: H.3.3 [Information Search and Retrieval]: Retrieval models

Keywords: adversarial retrieval, probability ranking principle

1. INTRODUCTION

The basic ad hoc document retrieval task is ranking documents in a corpus in response to a query by their relevance to the information need the query expresses. Numerous retrieval methods and frameworks have been devised throughout the years; e.g., the vector space model [28], the probabilistic approach, [29], the language modeling framework [23], the divergence from randomness framework [3], and learning-to-rank approaches [21].

The theoretical basis for all probabilistic retrieval methods, and under one interpretation or another, for all retrieval frameworks mentioned above, is the *probability ranking principle* (PRP) [26]. According to the PRP, documents in the corpus should be ranked by the probability that they are relevant to the query, so as to maximize the utility of the user who posted the query; relevance probabilities are assumed to be estimated using all information available to the search system [26][1].

There are two scenarios where the PRP may result in sub-optimal utility. The first is when users have different utility functions (i.e., personalization effects) [26]. The second scenario is when the relevance of one document depends on that of another; i.e., diversification effects [26, 7].

However, there is an additional fundamental aspect of the PRP that has not been accounted for in past work. The PRP, and retrieval methods devised based on the PRP, are based on a static view of the ad hoc retrieval setting. That is, the goal is to optimize user utility for the current query given the current *fixed* snapshot of the corpus[2]. The optimization is performed *regardless* of the incentives of the authors who created the documents, and thus does not account for potential post-retrieval effects — e.g., changes of the corpus, and therefore rankings for future queries, as a result of the ranking induced for the current query. More generally, an underlying implicit premise of the standard approach of designing ranking functions is that optimizing user utility along time can be based on providing the best possible relevance ranking for each query without accounting for future effects. We show in this paper that this static view falls short in dynamic and adversarial settings.

In adversarial retrieval settings such as the Web, documents might be changed by their authors, henceforth *publishers*, so as to have them promoted in rankings induced in response to queries. This practice is often referred to as search engine optimization [15]. Thus, the assumption that the ranking induced for one query will not affect the potential utility attained for future queries is simply wrong; that is, each induced ranking can lead to changes in the corpus upon which future search will be performed. Naturally, these changes can affect the utility attained for future queries. To illustrate this point, consider the following example. Suppose that a Web page is composed of a few sections, some of which contain valuable information which is unique to the page. In terms of page visibility, the publisher of the Web page is mainly concerned about how the page is ranked for queries that match the non-unique sections of the page. Having the page ranked low in response to these queries might

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.
ICTIR '15, September 27-30, 2015, Northampton, MA, USA.
©2015 ACM. ISBN 978-1-4503-3833-2/15/09$15.00.
DOI: http://dx.doi.org/10.1145/2808194.2809456 .

[1] There is also a definition of the PRP for interactive retrieval [13]. Interactive retrieval is outside the scope of this paper.
[2] In some cases, information about previous snapshots of the corpus might also be utilized [11].

lead the publisher to remove some of the unique sections, so as to better emphasize those which should "attract" the queries of concern. Such an action can lead to diminishing the utility of users looking for the unique information.

We present a novel game theoretic analysis of the adversarial ad hoc retrieval setting. The analysis accounts for the following facts (i) document publishers are "players" with *incentives*; namely, having their documents ranked high in response to some queries; (ii) the search engine is a *mediator* which affects the actions of publishers by the rankings it induces in response to queries; and therefore, (iii) the utility attained for the current query can affect that attained for future queries by the virtue of post-retrieval effects on the corpus (i.e., changes to documents). Our analysis assumes that all users have the same utility function (i.e., we do not account for personalization effects), and that the relevance of one document is independent of that of others (i.e., we do not account for the importance of diversifying search results but rather focus on their relevance). In that respect, our analysis is committed to the same conditions under which the PRP was shown to be optimal for a specific query and a static corpus.

A particular advantage of taking a game theoretic approach to modeling the adversarial retrieval setting, is the ability to make insightful statements about the steady state of the dynamics of the setting, namely, an equilibrium. The players — publishers in our case — who act selfishly do not have an incentive to deviate from the equilibrium. Thus, analyzing properties of the equilibrium is important for understanding the implications of a specific (dynamic) game. We note that any change to the actions that the players can take — in our case, the types of documents they can produce — and to the ranking function which serves as the mediator induces a new game with potentially new equilibria.

Inspired by standard practice in work on (algorithmic) game theory, our treatment of the adversarial retrieval setting as a game focuses on the notion of *social welfare*: the sum of utilities provided to users by rankings produced in response to their queries. Indeed, search engines should opt to maximize the social welfare of their users. More specifically, we analyze the worst possible social welfare attained in *any* equilibrium with respect to the best possible social welfare that can be attained. This concept, a.k.a. *price of anarchy* (PoA)[19, 27], is an important tool for analyzing games. In our case, the PoA is a means to contrasting different types of ranking functions in terms of the worst possible social welfare they can lead to in a steady state of the game.

We build up our analysis by starting from the case of having documents that discuss a single topic. We show that applying the PRP in this case, on a per-query basis, results in optimal social welfare if publishers are assumed to have the same writing quality for all topics. Later on we show that if this is not the case — i.e., publishers have differential writing qualities for topics — then applying the PRP is sub-optimal. As it turns out, the PRP is also sub-optimal in the case of having multi-topic documents and publishers with equal writing qualities for all topics. Interestingly, in several cases of our analysis, we show that introducing *randomization* into a ranking function can sometimes lead to social welfare that transcends that of applying the PRP.

The two main contributions of this paper are as follows:

- Presenting a novel game theoretic analysis of the adversarial ad hoc retrieval setting. The analysis accounts for the incentives of document authors to have their documents ranked high in response to queries.
- Showing that the standard practice of inducing a ranking based on the probability ranking principle — i.e., optimizing relevance ranking for a given query and static corpus — can be sub-optimal in terms of the overall utility attained by users of the search engine in the face of changes that the corpus goes through.

We believe that the sub-optimality of the PRP, and more generally, the sub-optimality entailed by the static view taken by current retrieval methods that are based on the PRP, have important implications. For example, our theoretical findings imply that learning ranking functions by optimizing rankings for a train set of queries and a static corpus, as is the current state-of-affairs [21], is a sub-optimal practice. That is, the learning should also account for post-retrieval effects in terms of potential changes applied to documents in the corpus by incentivized document authors.

2. RELATED WORK

The vast majority of published work on adversarial information retrieval has focused on spam classification (mainly content-based and link-based spam), and more generally, on estimating the quality and/or authoritativeness of Web pages (e.g., [12, 15, 4, 30, 20, 18, 8, 5, 9, 1, 2]). There has also been some recent work on estimating for which queries content manipulation of documents is more likely [24] and on estimating global term statistics in adversarial peer-to-peer search systems [25]. Work on adversarial classification treats the classifier and the data generator as two adversaries [10].

In contrast to all this prior work, we present a game theoretic modeling of the adversarial search (i.e., ad hoc retrieval) setting wherein publishers are incentivized to have their documents promoted in rankings. Our analysis does not depend on the nature of actions employed by publishers, namely, whether these are black hat (e.g, spamming), grey hat or white hat search engine optimization efforts [15]. Furthermore, showing that the probability ranking principle is sub-optimal in adversarial settings is a novel contribution of the work presented here.

There has been some (none game theoretic) work on modeling the past versions of a Web page — independently of those of other pages — so as to improve retrieval effectiveness [11]. In contrast, we provide a game theoretic analysis of the effect (specifically, in a steady state) of having multiple publishers change the content of their documents throughout time in response to multiple queries.

There is recent work on devising relevance-based dueling strategies for search engines [16]. Specifically, two search engines compete with each other assuming static publishers and documents. In contrast to our work, analysis of game theoretic aspects (e.g., equilibrium and social welfare) was not presented. Furthermore, our focus is on the competition between dynamic publishers (and documents) given that *some* search engine is used.

3. PRELIMINARIES

Let $[n]$ denote the set $\{1, 2, \ldots, n\}$, and $\Delta(S)$ denote the set of all probability distributions defined over a set S.

Given a distribution $d \in \Delta(S)$, we use $d(i)$ to denote the probability of event i ($\in S$) and $\mathrm{SUP}\,(d)$ to denote the set of all elements $i \in S$ for which $d(i) > 0$.

3.1 Game theory

We start with a standard definition of a game.

DEFINITION 1. *A strategic form n-player game is a tuple $G = (\{S_i\}_{i\in[n]}, \{U_i\}_{i\in[n]})$. For each player $i \in [n]$, S_i is the set of* **actions** *of the player. The set $S = S_1 \times S_2 \times \cdots \times S_n$ is called the set of* **pure strategy profiles**. *For each $i \in [n]$, $U_i : S \to \mathbb{R}$ is the* **utility function** *of player i.*

In a game, each player tries to maximize the (expected) value of his utility function. Note that the utility of a player depends on actions taken by all players, not only his own. The case of a 2-players game can be represented by a bi-matrix — a matrix in which each cell contains a 2-tuple: the first value is the payoff for the first player, and the second value is the payoff for the second player. Each row in the matrix represents an action of the first player, and each column represents an action of the second player; i.e., the matrix has $|S_1|$ rows and $|S_2|$ columns. For example, consider the following payoff matrix. According to the matrix,

	L	R
U	$(5, 2)$	$(0, 4)$
D	$(3, 6)$	$(1, 1)$

if the players chose actions D and L, the first (row) player payoff will be 3, while the second (column) player payoff will be 6. In a game, each player may choose an action randomly:

DEFINITION 2. *A* **mixed strategy** *for player i is a distribution $s_i \in \Delta(S_i)$.*

Using this definition, we extend the notion of a strategy profile so that $\Delta S \triangleq \Delta(S_1) \times \ldots \times \Delta(S_n)$ is the set of **mixed strategy profiles**.

We assume that each player i chooses his action according to $\Delta(S_i)$ *independently* of the actions selected by the other players. This allows a straight-forward generalization of the player's utility to the *expected* payoff over the joint distribution. That is, let $s = (s_1, \ldots, s_n) \in \Delta S$ be a mixed strategy profile, then

$$U_i(s) = \sum_{(a_1,\ldots,a_n)\in S} U_i(a_1, \ldots, a_n) \prod_{i=1}^{n} s_i(a_i).$$

We use $s_{-i} = (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n) \in \Delta S_{-i}$ to denote the strategies played by all players but i. A useful notion is the concept of **best response**. This is the set of actions a player may take to maximize his utility given the actions played by all other players:

DEFINITION 3. *Given a strategy profile s_{-i}, the best response of i is defined as $BR_i(s_{-i}) \triangleq argmax_{a_i\in S_i} U_i(a_i, s_{-i})$.*

Finally, this allows defining a "stable" state of a game: a strategy profile for which no player has an incentive to **deviate**, i.e., change his strategy. This concept is often referred to as the solution of a game.

DEFINITION 4. *A mixed strategy profile $s = (s_1, \ldots, s_n) \in \Delta S$ is a (mixed)* **Nash equilibrium**, *if for every player $i \in [n] : SUP(s_i) \subseteq BR_i(s_{-i})$. Equivalently, s is an equilibrium if $\forall i \in [n], \forall a'_i \in S_i : U_i(a'_i, s_{-i}) \le U_i(s_i, s_{-i})$. If s is a pure strategy profile, the equilibrium is called* pure.

	Football	Opera
Football	$(2, 1)$	$(0, 0)$
Opera	$(0, 0)$	$(1, 2)$

Table 1: Battle of the sexes payoff matrix.

For example, consider the two player game known as "battle of the sexes", where a couple has agreed to go on a date this evening, but forgot to agree on a venue. The husband prefers going to a football match, while the wife wants to see the opera, and they can not communicate. Both would prefer meeting at any place rather than not to meet at all. The game payoffs are given in Table 1. The game has two pure equilibria: jointly watching the match or attending the opera together. This game also has a mixed Nash equilibrium, where the husband goes to the match with probability 2/3 while the wife attends the opera with the same probability. Note that if each were to pick a venue uniformly, they would have met with a higher probability. However, the different preferences of the two incentivize each to try his/hers place of choice with a higher probability. A fundamental result in game theory is Nash Theorem:

THEOREM 1. *[22] Every game with a finite number of players and action sets has a mixed strategies equilibrium.*

A common metric for measuring the "quality" of a given mixed strategy profile is called **social welfare**. The social welfare is defined as the sum of (expected) utilities received by the players playing it; that is, for $s \in \Delta S$ we define $SW(s) \triangleq \sum_{i\in[n]} U_i(s)$.

3.2 Price of anarchy

Given a game, an interesting question is "how inefficient" the system could get in a stable state, due to the selfish behavior of the players. Specifically, an equilibrium of the game is compared with an idealized scenario where all players selflessly collaborate to maximize the profit of the group which is measured by social welfare. A popular measure of this inefficiency is called the **Price of Anarchy** (PoA) [19, 27]. Formally, the price of anarchy is the ratio between the sum of utilities achieved by an optimal strategy profile (i.e., $OPT = \max_{s\in\Delta S} SW(s)$), and the social welfare of the *worst* equilibrium, i.e., the one with the lowest social welfare. Going back to the example from Table 1, when both are going together to either the opera or the football match, we have an optimal social welfare of 3, while the mixed equilibrium only provides a social welfare of 4/3. Thus, the price of anarchy for the game is $\frac{3}{4/3} = 2.25$.

4. ADVERSARIAL AD HOC RETRIEVAL

As already noted, in adversarial retrieval settings such as the Web, the goal of many owners of documents, henceforth referred to as *publishers*, is to have their documents ranked as high as possible in response to queries of interest. Specifically, a ranking induced for a query by a search engine might incentivize publishers of documents that are not highly ranked to change their documents for future rankings. Naturally, this is an on-going process.

We model the dynamics just mentioned as a game. The publishers are players whose actions are document production; i.e., a publisher decides which document to produce at any point in time. The search engine plays the role of a *mediator*. The choice of a ranking function affects the publishers' utilities by the virtue of the induced rankings.

We make the assumption that publishers' utilities and users' utilities are aligned. That is, the payoff a publisher gets for a document with respect to a user query reflects the utility the user attains from reading the document in response to the query. In other words, the premise is that user satisfaction will lead to publisher satisfaction. In that respect, the social welfare of publishers (as players) attained in a game is considered to be the social welfare of the search engine's users. Analysis of more complicated situations, wherein the publishers' and users' utilities are not aligned, is left for future work.

Our goal here is to study the optimality, or lack thereof, of ranking functions that obey the probability ranking principle (PRP) in adversarial retrieval settings; that is, functions that rank documents in response to a query by the relevance (probability) of the documents. To that end, we describe a few games that are defined by the choice of the ranking function and the assumptions about the type of documents the publishers can produce. We assume that in each game the ranking function we want to analyze is fixed. Accounting for changes of the ranking function throughout the game is outside the scope of this paper.

To analyze a ranking function, we examine the equilibria of the game it entails. Specifically, we focus on the price of anarchy of the game (PoA; see Section 3.2), as is common in work on algorithmic game theory. Indeed, this practice allows to reason about the worst case steady state of the game (equilibrium), in terms of the social welfare attained, given the selfish behavior of players — publishers in our case.

4.1 Single-topic documents

Here and after, we assume that each query is about a single topic. We start our analysis in this section with a basic model where every document is also assumed to be about a single topic. The model assumes, in addition, that all publishers have equal writing quality for all topics and that each publisher may write on a topic of his choice. We show that in this model, the probability ranking principle (PRP) is optimal; i.e., using a ranking function that obeys the PRP results in optimal price of anarchy. In Sections 4.2 and 4.3 we extend this model by allowing multi topic documents, and differential writing qualities, respectively, and show that the PRP is no longer optimal in these cases.

4.1.1 Definitions

A **Publishers Game** $PG = (n, m, \mathcal{D}; \mathcal{R})$ is an n **publishers** game over m **topics**. Each publisher selects a topic to write on; i.e., the possible actions are topics. In addition, **queries** are posted, each of which on a specific topic, according to a public query distribution $\mathcal{D} \in \Delta([m])$. Throughout the paper, we assume without loss of generality that $\mathcal{D}(1) \geq \mathcal{D}(2) \geq \ldots \geq \mathcal{D}(m)$. That is, the first topic is always the most frequently queried, then the second, etc.

For the sake of simplifying the analysis, we assume that the user utility with respect to his query is determined solely based on the relevance of the first ranked document. While in practice users examine also documents in lower ranks, it is a well known fact that in the Web setting users pay most of their attention to the most highly ranked documents [17]. In fact, users often prefer to reformulate their queries rather than spend additional time examining lower ranked documents. Furthermore, it was shown that the precision at rank 1 (i.e., 0 or 1) is highly correlated (with Pearson correlation above 0.5) with average precision at cutoff 1000, which is a standard measure for quantifying the effectiveness of a retrieved list [6]. Thus, we assume that the utility the user attains from the top-most document is highly correlated with the utility the user will attain by examining additional documents further down the ranked list.

Given a query and the set of topics selected by the publishers, the *Ranking Function* $\mathcal{R} : [m] \times [m]^{[n]} \rightarrow \Delta([n])$, selects a probability distribution over the documents (i.e., over the players) using which the documents (players) will be ranked; specifically, the probability assigned to a document corresponds to the probability that it will be ranked first. We note that this definition is different than that of standard ranking functions in two respects. First, it does not specify how to rank documents for ranks two and below. Second, the ranking function is allowed to be *probabilistic*, randomizing over the set of documents in the corpus. In contrast, ranking functions used in the information retrieval literature are deterministic. While at first glance randomization may not seem to provide additional value, we show that it could actually result in improved social welfare even if it hurts retrieval effectiveness for a specific query.

As an example of the notation used for ranking functions, consider $\mathcal{R}(1, \langle 1, 1, 2 \rangle) = \langle 0.4, 0.4, 0.2 \rangle$. Given a query on the first topic, if players 1 and 2 wrote on topic 1, and player 3 created a document on topic 2, then each of players 1 and 2 will be ranked first with probability 0.4, while player 3 will be ranked first with probability 0.2. We use $\mathcal{R}_i : [m] \times [m]^{[n]} \rightarrow [0, 1]$ to denote the probability assigned to player i for the given input; in our example, $\mathcal{R}_3(1, \langle 1, 1, 2 \rangle) = 0.2$.

We define the payoff for publisher i as 1 for queries he is ranked first on, and 0 for all other queries. His expected payoff is then $U_i(a_i, a_{-i}) = \mathcal{D}(a_i) \cdot \mathcal{R}_i(a_i, (a_i, a_{-i}))$; i.e., the probability that a query on his selected topic arrives times the probability he will be ranked first.

We define the probability ranking principle ranking, **PRP ranking** in short, denoted R_{PRP}, to be the function which always ranks first the document most relevant to a query. Indeed, the PRP [26] states that documents should be ranked by their relevance probabilities, and we care here about the highest ranked document. In this model, the PRP translates to choosing uniformly a document written on the query topic for the first rank. For example, for the two publishers game $(2, m, \mathcal{D}; R_{PRP})$ where the PRP serves for ranking, the payoffs can be expressed by the matrix in Table 2.

	1	2	…	m
1	$\left(\frac{\mathcal{D}(1)}{2}, \frac{\mathcal{D}(1)}{2}\right)$	$(\mathcal{D}(1), \mathcal{D}(2))$		$(\mathcal{D}(1), \mathcal{D}(m))$
2	$(\mathcal{D}(2), \mathcal{D}(1))$	$\left(\frac{\mathcal{D}(2)}{2}, \frac{\mathcal{D}(2)}{2}\right)$		$(\mathcal{D}(2), \mathcal{D}(m))$
⋮			⋱	
m	$(\mathcal{D}(m), \mathcal{D}(1))$	$(\mathcal{D}(m), \mathcal{D}(2))$		$\left(\frac{\mathcal{D}(m)}{2}, \frac{\mathcal{D}(m)}{2}\right)$

Table 2: Payoff matrix for the game $(2, m, \mathcal{D}; R_{PRP})$.

4.1.2 Price of anarchy

We focus on the two players game, where a couple of rival publishers compete for the same incoming queries flow, each trying to maximize his payoff by being ranked first. Our goal in this section is showing that in this basic model, the PRP remains optimal. We start with a theorem showing that the PRP ranking yields a game, in which every equilibrium is at most 1.5 times worse than the optimal scenario in terms of social welfare. The optimal scenario is having each of the two publishers write about a different topic among the two most frequently queried topics.

THEOREM 2. *Let $\mathcal{D} \in \Delta([m])$ be a query distribution and let $PG = (2, m, \mathcal{D}; R_{PRP})$ be a two player Publishers Game, then the price of anarchy of PG is at most* 1.5.

The proof, presented in Appendix A, is based on analysis of two cases: when publishers write deterministically on a single topic, and when they choose to randomize over several topics. We show that in any case, the worst equilibrium is reached when both publishers choose a topic only from the two most frequently queried topics.

We complement Theorem 2 by proving that no ranking function is able to achieve price of anarchy smaller than 1.5, for any query distribution. Together, the theorems show that the PRP is optimal in the case where each document (and query) is about a single topic, and all publishers have the same writing quality for all topics. As noted above, optimality refers to the price of anarchy unless otherwise noted.

THEOREM 3. *For any ranking function $\mathcal{R}$, there exists a distribution $\mathcal{D} \in \Delta([m])$, such that the Publishers Game $PG = (2, m, \mathcal{D}; \mathcal{R})$ has a price of anarchy of* at least 1.5.

Intuitively, we show that for the query distribution in which the first topic is queried two-thirds of the time, while the second is queried otherwise, $\mathcal{R}$ can not decrease the price of anarchy below 1.5. The full proof is in Appendix B.

For the case of $m = 2$ topics, we strengthen the result of Theorem 3 by showing that for *any* distribution, one can not improve the social welfare obtained by R_{PRP} using a different ranking function.

THEOREM 4. *For any distribution $\mathcal{D} \in \Delta([2])$ and ranking function $\mathcal{R}$, the social welfare of $(2, 2, \mathcal{D}; \mathcal{R})$ is at most the social welfare of $(2, 2, \mathcal{D}; R_{PRP})$.*

The proof of the theorem is provided in Appendix C. Thus, we get that in the two topics setting, the PRP ranking is optimal in a stronger sense than that considered above (i.e., price of anarchy): there exists *no* query distribution for which an alternative function could achieve higher social welfare than that attained by the PRP.

While we showed that the PRP is optimal in this model, it is not the case in the models considered next. Section 4.2 discusses a setting in which publishers may produce multi-topic documents. Section 4.3 presents a model which accounts for publishers with differential writing qualities.

4.2 Multi-topic documents

We next discuss an extension of the Publishers Game that allows publishers to compose documents about multiple topics. Formally, a **Multiple Topics Publishers Game** $MTPG = (n, m, c, \mathcal{D}; \mathcal{R})$ is similar to Publishers Game, except that publishers may now choose an action from $[m] \cup \binom{[m]}{2}$; i.e., either writing a document on a single topic or on two topics. We assume that a multiple topics document has equal proportions of the two topics. When a publisher writes on two topics, an additional parameter, $c \in [0, 1]$, is the payoff he receives when ranked first on either of these topics. Assuming that documents are of equal length, the PRP ranking amounts here to determining relevance level (grade) based on the prevalence of relevance information in the document. This graded relevance definition was also applied in work on focused retrieval [14]. Specifically, a document dedicated entirely to the query topic is preferred to any document that discusses two topics. For two documents that discuss two topics, the one which discusses the query topic is preferred to the one which does not; if both documents discuss the query topic, one of the two is randomly selected to be ranked first.

We focus on the special case of two documents and two publishers, showing that even in this restricted scenario, the PRP is not always optimal. We start by showing that if the multi-topic document payoff c is small, then the PRP ranking remains optimal.

THEOREM 5. *Let $c \in [0, 1/2]$. For any distribution $\mathcal{D} \in \Delta([m])$, the Multiple Topics Publishers Game $MTPG = (2, m, c, \mathcal{D}; R_{PRP})$ achieves a price of anarchy of at most* 1.5.

The idea behind the proof, presented in Appendix D, is showing that when multi-topic documents provide just a small benefit, the publishers will never write such documents; in this case, we can use Theorem 3 for the price of anarchy bound.

The next theorem shows that for such small c values, the PRP is indeed optimal as any ranking function will incur a price of anarchy of 1.5 for some query distribution.

THEOREM 6. *Let $c \in [0, 1/2]$. There exists a distribution $\mathcal{D} \in \Delta([m])$, such that for* any *ranking function $\mathcal{R}$, the Multiple Topics Publishers Game $MTPG = (2, m, c, \mathcal{D}; \mathcal{R})$ has a price of anarchy of* at least 1.5.

Similarly to Theorem 3, we consider a query distribution where the first topic is queried two-thirds of the time, and otherwise the second topic is queried. We show that for this distribution, no ranking function could reduce the price of anarchy when c is small.

Next, when c is sufficiently large, we show an example in which the anarchy entailed by using the PRP ranking could be completely avoided by using an alternative ranking function which does not adhere to the PRP; that is, a social welfare that approaches the optimal can be attained.

EXAMPLE 1. *Consider the distribution $\mathcal{D}(1) = 2/3$, $\mathcal{D}(2) = 1/3$. For any c, the price of anarchy of $MTPG_{PRP} = (2, 2, c, \mathcal{D}; R_{PRP})$ is 1.5, while there exists a ranking function, R_{prob}, such that the price of anarchy of $MTPG_{prob} = (2, 2, c, \mathcal{D}; R_{prob})$ could be arbitrarily close to 1 as $c \to 1$.*

For the example, we propose a non-deterministic, probabilistic, ranking function[3] R_{prob}. The function randomizes the relative ranking of a multi-topic document and a pure

[3]As explained in Section 4.1, by probabilistic ranking function, we mean a *non-deterministic* ranking function that applies randomization. This should be differentiated from probabilistic retrieval methods that are deterministic and rank documents by relevance probabilities [29].

document written on the first topic, whenever a query on the first topic is posted. The complete proof is provided in Appendix E.

We conclude this section by showing a large family of games in which the PRP is not optimal, and propose a probabilistic (randomized) ranking function instead. This finding implies that introducing randomness to the *ranking* of documents may result in a more content enriched corpus in the long run, by providing incentives for publishers to diversify the topics they write about.

THEOREM 7. *Let* $c \in (1/2, 1]$ *,* $p \in [2/3, \frac{2c}{c+1})$*. Consider the distribution* $\mathcal{D}(1) = p, \mathcal{D}(2) = 1-p$*. There exists a ranking function* R_{prob}*, such that the price of anarchy of* $(2, 2, c, \mathcal{D}; R_{PRP})$ *is strictly larger than that of* $(2, 2, c, \mathcal{D}; R_{prob})$ *.*

The core idea underlying the proof of the theorem, which can be found in Appendix D, is to carefully construct a probabilistic ranking function which occasionally ranks the multi-topic document higher than a document written entirely on the query topic.

4.3 Publishers with differential writing qualities

Heretofore, we assumed that every publisher can write about every topic with the same quality. We now turn to analyze a setting wherein this is not the case. Specifically, we explore an extension of the Publishers Game to a setting where publishers write on a single topic, but have different writing qualities. Formally, a **Different Qualities Publishers Game** $DQPG = (n, m, \mathcal{Q}; \mathcal{R})$ is similar to the Publishers Game, except that publishers now have individual payoff functions representing the utility of the document they create for a queried topic. This setting amounts to assuming, for example, that the utility (for the user, and consequently for the publisher, per our working assumption) attained from the same document produced by two different publishers is different and depends on their writing quality for the topic. Indeed, the less the publisher is authoritative about the topic, the more effort the user will have to put into verifying the content (e.g., by consulting other sources of information), thereby decreasing his utility.

The DQPG game has a **Quality Matrix** parameter, $\mathcal{Q} \in [0,1]^{[n]\times[m]}$, where $\mathcal{Q}_{i,j}$ is the payoff for publisher i if he is ranked first on topic j. Note that the query distribution parameter has been omitted as we may normalize the Quality Matrix to factor the topic frequencies. For example, if the first topic is queried $2/3$ of the time, and otherwise the second topic is queried, then we multiply the first column of $\mathcal{Q}$ by $2/3$ and the second column by $1/3$. In this setting, we assume that in the PRP ranking, R_{PRP}, the *highest quality* document written on the queried topic is ranked first. In case multiple documents have the same quality, it uniformly selects a top-quality document. That is, originally, the PRP was shown to optimize utility by *implicitly* assuming equal writing qualities of publishers; thus, relevance amounted to utility for a given query. Here, we stick to the original goal of the PRP, optimizing utility, and account for the fact that relevance is not necessarily coupled with utility.

First, we show a tight bound for the price of anarchy achieved by the PRP ranking.

THEOREM 8. *For any Quality Matrix* $\mathcal{Q} \in [0,1]^{[n]\times[m]}$*, such that players have distinct writing qualities for any given topic, the Different Qualities Publishers Game* $DQPG = (n, m, \mathcal{Q}; R_{PRP})$ *has a price of anarchy of at most* 2*.*

The essence of the proof, provided in Appendix G, is reducing the problem to finding a matching in the bipartite graph which has the publishers on one partite and the topics on the other. We show that any equilibrium in the game corresponds to a matching which may be generated by the greedy algorithm and that the approximation rate of the algorithm provides a bound on the price of anarchy.

Complementing the theorem, we show that the price of anarchy analysis for R_{PRP} is tight.

THEOREM 9. *For any* $\epsilon > 0$*, there exists a matrix* $\mathcal{Q}^{\epsilon}$*, such that the price of anarchy of* $(n, m, \mathcal{Q}^{\epsilon}; R_{PRP})$ *is* $2 - \epsilon$*.*

The result is obtained for the simple case where one publisher has writing quality for multiple topics, while another publisher can write mainly on a specific topic. If the first player has even a slight writing quality advantage over the second on that specific topic, he may prefer writing about it, and no publisher would satisfy the information need for the remaining topics. The proof is found in Appendix H.

It is important to note that if a ranking function could act arbitrarily, a search engine could abuse this model to force the optimal strategy, which is the *maximum weight matching* of the complete bipartite graph over $[n] \cup [m]$, where the weight of an edge (i, j) is $\mathcal{Q}_{i,j}$. For example, consider the ranking function which given a query on topic j, ranks first the publisher i, if (i, j) is a part of the maximum weight matching, and i wrote a document on topic j. Otherwise, it places a completely irrelevant document at the top position. In this case, every publisher realizes that the only way users with information need relevant to his page reach his document is by cooperation with the engine, writing on the topic assigned to him in the matching.

In order to avoid ranking functions which dictate for each player the topic to write on, we restrict the discussion to **Fair Ranking Functions**. A Fair Ranking Function is a function which given a query, assigns first-rank probabilities (i.e., probability to be ranked first) to documents written on the topic, without considering the writing qualities of the publishers on *other* topics. More formally, a Fair Ranking Function is a set of functions, $\mathcal{R} = \{\mathcal{R}^k : [0,1]^{[k]} \to \Delta([k])\}_{k\in[n]}$, such that given a query on a topic that k publishers wrote about, $\mathcal{R}^k$ is applied on the qualities of these publishers to determine the ranking, and the ranker is not allowed to consider the writing qualities of publishers on topics that they did not write on.

Theorem 9 shows that for *any* values of n, m (i.e., number of publishers and number of topics, respectively) the price of anarchy of the PRP is 2. We complement this by showing that for the case of 2-player games, R_{PRP} is not optimal, even when compared only with Fair Ranking Functions.

THEOREM 10. *There exists a Fair Ranking Function* R_{uni}*, such that for any Quality Matrix* $\mathcal{Q} \in [0,1]^{[2]\times[m]}$*, the Different Qualities Publishers Game* $DQPG = (2, m, \mathcal{Q}; R_{uni})$ *achieves a price of anarchy of at most* 1.7*.*

To prove the theorem, we consider a probabilistic ranking function R_{uni} which makes a fair coin flip to select the top ranked document, whenever the difference in quality of given documents is "small enough". A more detailed proof sketch for the $n = m = 2$ case appears in Appendix I. The complete proof, and its extension to general number of topics m are omitted due to lack of space.

We conclude the analysis by showing a lower bound on the price of anarchy achievable by *any* Fair Ranking Function.

THEOREM 11. *For any Fair Ranking Function $\mathcal{R}$ and number of topics m, there exists a Quality Matrix $\mathcal{Q} \in [0,1]^{[2]\times[m]}$, such that the price of anarchy of the Different Qualities Publishers Game $DQPG = (2, m, \mathcal{Q}; R_{uni})$ is at least* 1.5.

In Appendix J, we present a proof which constructs two games with different quality matrices. We show that any ranking function achieves social welfare which is at most two thirds of the optimum, at least for one of the games.

5. CONCLUSIONS

We presented a novel game theoretic analysis of the ad hoc document retrieval task in adversarial settings. The analysis accounts for the incentives of authors to have their documents ranked high in response to queries. Thus, the analysis provides formal grounds for modeling the dynamic nature of the adversarial retrieval setting that results from authors consistently changing their documents so as to promote them in rankings. We performed the analysis for different types of documents (namely, single topic versus multi topic) and by using different assumptions about the writing qualities of authors. One of our most important findings is that the probability ranking principle (PRP), which is the theoretical foundation of most ad hoc retrieval methods, can be sub-optimal in the adversarial setting. Specifically, we showed that in some cases, introducing randomization into the document ranking function can result in user utility higher than that attained by applying the PRP.

Acknowledgements

We thank the reviewers for their comments. This work was supported by and carried out at the Technion-Microsoft Electronic Commerce Research Center.

6. REFERENCES

[1] *AIRWeb — International Workshop on Adversarial Information Retrieval on the Web*, 2005–2009.

[2] *WICOW/AIRWeb Workshop on Web Quality (WebQuality)*, 2012.

[3] G. Amati and C. J. van Rijsbergen. Probabilistic models of information retrieval based on measuring the divergence from randomness. *ACM Transactions on Information Systems*, 20(4):357–389, 2002.

[4] A. A. Benczúr, K. Csalogány, T. Sarlós, and M. Uher. Spamrank – fully automatic link spam detection. In *AIRWeb 2005, First International Workshop on Adversarial Information Retrieval on the Web, co-located with the WWW conference, Chiba, Japan, May 2005*, pages 25–38, 2005.

[5] M. Bendersky, W. B. Croft, and Y. Diao. Quality-biased ranking of web documents. In *Proc. of WSDM*, pages 95–104, 2011.

[6] O. Butman, A. Shtok, O. Kurland, and D. Carmel. Query-performance prediction using minimal relevance feedback. In *Proceedings of ICTIR*, page 7, 2013.

[7] J. G. Carbonell and J. Goldstein. The use of MMR, diversity-based reranking for reordering documents and producing summaries. In *Proc. of SIGIR*, pages 335–336, 1998.

[8] G. V. Cormack. TREC 2007 spam track overview. In *Proc. of TREC*, 2007.

[9] G. V. Cormack, M. D. Smucker, and C. L. A. Clarke. Efficient and effective spam filtering and re-ranking for large web datasets. *Information Retrieval*, 14(5):441–465, 2011.

[10] N. N. Dalvi, P. Domingos, Mausam, S. K. Sanghai, and D. Verma. Adversarial classification. In *Proc. of KDD*, pages 99–108, 2004.

[11] J. L. Elsas and S. T. Dumais. Leveraging temporal dynamics of document content in relevance ranking. In *Proc. of WSDM*, pages 1–10, 2010.

[12] D. Fetterly, M. Manasse, and M. Najork. Spam, damn spam, and statistics: Using statistical analysis to locate spam web pages. In *Proceedings of WebDB*, pages 1–6, 2004.

[13] N. Fuhr. A probability ranking principle for interactive information retrieval. *Information Retrieval*, 11(3):251–265, 2008.

[14] S. Geva, J. Kamps, and R. Schenkel, editors. *Focused Retrieval of Content and Structure, 10th International Workshop of the Initiative for the Evaluation of XML Retrieval, INEX 2011, Saarbrücken, Germany, December 12-14, 2011, Revised Selected Papers*, volume 7424 of *Lecture Notes in Computer Science*. Springer, 2012.

[15] Z. Gyöngyi and H. Garcia-Molina. Web spam taxonomy. In *Proc. of AIRWeb 2005*, pages 39–47, 2005.

[16] P. Izsak, F. Raiber, O. Kurland, and M. Tennenholtz. The search duel: a response to a strong ranker. In S. Geva, A. Trotman, P. Bruza, C. L. A. Clarke, and K. Järvelin, editors, *Proc. of SIGIR*, pages 919–922. ACM, 2014.

[17] T. Joachims, L. A. Granka, B. Pan, H. Hembrooke, and G. Gay. Accurately interpreting clickthrough data as implicit feedback. In R. A. Baeza-Yates, N. Ziviani, G. Marchionini, A. Moffat, and J. Tait, editors, *Proc. of SIGIR*, pages 154–161, 2005.

[18] T. Jones, D. Hawking, and R. Sankaranarayana. A framework for measuring the impact of web spam. In *Proc. of ADCS2007*.

[19] E. Koutsoupias and C. Papadimitriou. Worst-Case Equilibria. In *Proc. of STACS*, 1999.

[20] V. Krishnan and R. Raj. Web spam detection with Anti-Trust rank. In *Proceedings of AIRWeb*, pages 37–40, 2006.

[21] T.-Y. Liu. Learning to rank for information retrieval. *Foundations and Trends in Information Retrieval*, 3(3), 2009.

[22] J. F. Nash. Equilibrium points in n-person games. *Proc. of the National Academy of Sciences of the United States of America*, 36.1:48–49, 1950.

[23] J. M. Ponte and W. B. Croft. A language modeling approach to information retrieval. In *Proc. of SIGIR*, pages 275–281, 1998.

[24] F. Raiber, K. Collins-Thompson, and O. Kurland. Shame to be sham: Addressing content-based grey hat search engine optimization. In *Proc. of SIGIR*, pages 1013–1016, 2013.

[25] S. Richardson and I. J. Cox. Estimating global statistics for unstructured P2P search in the presence of adversarial peers. In *Proc. of SIGIR*, pages 203–212, 2014.

[26] S. E. Robertson. The probability ranking principle in IR. *Journal of Documentation*, pages 294–304, 1977.

[27] T. Roughgarden and E. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, April 2002.

[28] J. Salton, A. Wong, and C. S. Yang. A vector space model for automatic indexing. *Communications of the ACM*, 18(11):613–620, 1975.

[29] K. Sparck Jones, S. Walker, and S. E. Robertson. A probabilistic model of information retrieval: development and comparative experiments - part 1. *Information Processing and Management*, 36(6):779–808, 2000.

[30] B. Wu and B. D. Davison. Identifying link farm spam pages. In *Proc. of WWW*, pages 820–829, 2005.

APPENDIX

A. PROOF OF THEOREM 2

PROOF. The optimal scenario we compare to, is when some publisher writes about topic 1, while the other composes a document on topic 2. The optimal social welfare is therefore $OPT = \mathcal{D}(1) + \mathcal{D}(2) \leq 1$. Next, we examine a few possible equilibria:

1. Each player writes on a different topic.
   When such equilibrium exists, players must write about topics 1 and 2 (or other topics with the same frequency, if such exist). The social welfare of the equilibrium is then $\mathcal{D}(1)+\mathcal{D}(2)$, and the players reach optimal welfare.

2. Both players write about topic 1 with probability 1.
   In this scenario, $\mathcal{D}(1) \geq 2/3$, as the publishers have no incentive to deviate to topic 2. This means the social welfare of the equilibrium is $\mathcal{D}(1)$, and the price of anarchy is at most $\frac{1}{2/3} = 1.5$.

3. The players play a symmetric equilibrium (i.e., both have the same mixed strategy), randomizing their actions over the first two topics.
   As both players have topic 2 in their support, we know that $\mathcal{D}(2) \geq \mathcal{D}(1)/2$. Assume that each player writes about the first topic with probability p. The payoff of writing about topic 1, assuming that the other publishers follows this strategy, is $\mathcal{D}(1)/2$ with probability p (as the ranking function would rank each of the publishers randomly), or $\mathcal{D}(1)$ if the other player chose to write about topic 2, which happens with probability $1-p$. Therefore the utility is $U(1) = p \cdot \mathcal{D}(1)/2 + (1-p) \cdot \mathcal{D}(1)$. Similarly, the utility of writing about the second topic is $U(2) = p \cdot \mathcal{D}(2) + (1-p) \cdot \mathcal{D}(2)/2$. As both topics are in the support of the players' strategy, the utility of writing on each of them must be equal, i.e. $U(1) = U(2)$. This give us $p = \frac{2\mathcal{D}(1)-\mathcal{D}(2)}{\mathcal{D}(1)+\mathcal{D}(2)}$. Since both publishers play the same mixed strategy, the social welfare of the equilibrium is $P = 2 \cdot U(1) = 2 \cdot U(2) = \frac{3\mathcal{D}(1)\mathcal{D}(2)}{\mathcal{D}(1)+\mathcal{D}(2)}$. Finally, we can express the price of anarchy as

$$\frac{OPT}{P} = \frac{\mathcal{D}(1)+\mathcal{D}(2)}{\frac{3\mathcal{D}(1)\mathcal{D}(2)}{\mathcal{D}(1)+\mathcal{D}(2)}} = \frac{(\mathcal{D}(1)+\mathcal{D}(2))^2}{3\mathcal{D}(1)\mathcal{D}(2)}.$$

   In the domain $0 \leq \mathcal{D}(1)/2 \leq \mathcal{D}(2) \leq \mathcal{D}(1) \leq 1-\mathcal{D}(2)$, this function is bounded by $3/2$, which is reached on the line $\mathcal{D}(1) = 2\mathcal{D}(2)$.

4. One of the players does not have the second topic in the support of his strategy.
   Without loss of generality, we assume that player 2 is the one which never writes about topic 2. In this case, since we already covered all equilibria in which players restrict themselves only to the first two topics, the second player must have the third topic in his support (or a different topic with the same query frequency). The first player then must earn at least $\mathcal{D}(2)$, as he would be the only publisher to write on the second topic, should he choose to do so. Player 2 may randomize between the first and third topic, but since he is the only publisher to (occasionally) write on topic 3, his utility is $\mathcal{D}(3)$. Also, since the players can always earn $\mathcal{D}(1)/2$ by writing on the first topic, we get that $\mathcal{D}(2), \mathcal{D}(3) \geq \mathcal{D}(1)/2$. Putting it together, we get a social welfare of $\mathcal{D}(2) + \mathcal{D}(3)$, which means price of anarchy of

$$\begin{aligned}\frac{OPT}{P} &\leq \frac{\mathcal{D}(1)+\mathcal{D}(2)}{\mathcal{D}(2)+\mathcal{D}(3)} \leq \frac{\mathcal{D}(1)+\mathcal{D}(2)}{\mathcal{D}(2)+\mathcal{D}(1)/2} \\ &= \frac{\mathcal{D}(1)+\mathcal{D}(2)}{2\mathcal{D}(2)/3+\mathcal{D}(2)/3+\mathcal{D}(1)/2} \\ &\leq \frac{\mathcal{D}(1)+\mathcal{D}(2)}{2\mathcal{D}(2)/3+\mathcal{D}(1)/6+\mathcal{D}(1)/2} = \frac{3}{2}\end{aligned}$$

5. Both players are playing an equilibrium in which the support of each contains topics 1 and 2.
   In this case, we argue that each player must earn at least as much as he did in equilibrium in which both players only write on the first two topics (case 3). Let $p = \frac{2\mathcal{D}(1)-\mathcal{D}(2)}{\mathcal{D}(1)+\mathcal{D}(2)}$ be the 2-topics symmetric equilibrium probability computed in 3. Notice that for an equilibrium $< s_1, s_2 >$, we get that the utilities received by the players are:

$$U_1 = \mathcal{D}(1)(1 - s_2(1)/2) = \mathcal{D}(2)(1 - s_2(2)/2)$$

$$U_2 = \mathcal{D}(1)(1 - s_1(1)/2) = \mathcal{D}(2)(1 - s_1(2)/2)$$

   Now if $\{1,2\} \subsetneq SUP(s_i)$ then either $s_i(1) < p$ or $s_i(2) < 1-p$. This means that the profit for the *other* player has to increase. Since the arguments holds for both players, we get that each gains at least $\frac{3\mathcal{D}(1)\mathcal{D}(2)}{2(\mathcal{D}(1)+\mathcal{D}(2))}$ and the bound from 3 give us the required $\frac{3}{2}$ bound on the price of anarchy.

Since we covered all possible equilibria, we established that the price of anarchy is 1.5. □

B. PROOF OF THEOREM 3

PROOF. Consider the simple scenario where queries about the first topic arrive with probability 2/3, while the second topic is queried 1/3 of the time; i.e., $\mathcal{D}(1) = 2/3, \mathcal{D}(2) = 1/3, \forall i \in [m] \setminus \{1,2\} : \mathcal{D}(i) = 0$. In the PRP-based game, $PG = (2, m, \mathcal{D}; R_{PRP})$, the players' payoffs are given by

(1/3, 1/3)	(2/3, 1/3)
(1/3, 2/3)	(1/6, 1/6)

. The ranking function $\mathcal{R}$ is allowed to choose the first-rank probabilities when the players write on different topics; i.e., the probability of ranking a non-relevant document first, rather than one written on the queried topic. This means that $\mathcal{R}$ may modify the game to

(1/3, 1/3)	(α, β)
(β, α)	(1/6, 1/6)

, for any values of α, β satisfying $0 \leq \alpha \leq 2/3, 0 \leq \beta \leq 1/3$. In the new game, regardless of the values of α and β, there exists an equilibrium in which both players write on topic 1. In this equilibrium, each player payoff is 1/3, while the utility of deviating to the second topic is $\beta \leq 1/3$. This means that the price of anarchy of $\mathcal{R}$ is at least $\frac{1}{2/3} = 1.5$. □

C. PROOF OF THEOREM 4

PROOF. When R_{PRP} is used, the game is a special case of the one presented in Table 2, and is given by the matrix

$\left(\frac{\mathcal{D}(1)}{2}, \frac{\mathcal{D}(1)}{2}\right)$	$(\mathcal{D}(1), \mathcal{D}(2))$
$(\mathcal{D}(2), \mathcal{D}(1))$	$\left(\frac{\mathcal{D}(2)}{2}, \frac{\mathcal{D}(2)}{2}\right)$

. If $\mathcal{D}(1) \geq 2/3$, the game has a unique equilibrium in which both players write about the first topic. This means that the social welfare of the game will be $\mathcal{D}(1)$. When $\mathcal{D}(1) \in [1/2, 2/3)$, the game has an optimal social welfare equilibrium, in which publishers write on different topics. The game also has a mixed strategies equilibrium, in which each writes about the first topic with probability $p = 3\mathcal{D}(1) - 1$. This equilibrium has a social welfare of $3\mathcal{D}(1)\mathcal{D}(2)$. Similarly to Theorem 3, the ranking function $\mathcal{R}$ may rank a document on a different topic before a relevant topic, hence it may modify the game to

$\left(\frac{\mathcal{D}(1)}{2}, \frac{\mathcal{D}(1)}{2}\right)$	(α, β)
(β, α)	$\left(\frac{\mathcal{D}(2)}{2}, \frac{\mathcal{D}(2)}{2}\right)$

for any values of α, β such that $0 \leq \alpha \leq \mathcal{D}(1), 0 \leq \beta \leq \mathcal{D}(2)$. In the case where $\alpha \geq \frac{\mathcal{D}(2)}{2}$ and $\beta \geq \frac{\mathcal{D}(1)}{2}$, this game has a symmetric mixed strategy equilibrium where each player writes about the first topic with probability $p \triangleq \frac{2\alpha+\mathcal{D}(1)-1}{2\alpha+2\beta-1}$. In all other cases, the game has only pure-strategy equilibrium where both surely write about topic 1. The pure strategy equilibrium gives a social welfare of $\mathcal{D}(1)$. This is strictly lower than the $3\mathcal{D}(1)\mathcal{D}(2)$ achieved by R_{PRP} in the $\mathcal{D}(1) \in [1/2, 2/3)$ case, and identical to it when $\mathcal{D}(1) \geq \frac{2}{3}$. The mixed strategies equilibrium gives a social welfare of $\mathcal{D}(1) \cdot p + (\alpha+\beta) \cdot (1-p) = \frac{4\alpha\beta - \mathcal{D}(1)\mathcal{D}(2)}{2\alpha+2\beta-1}$. Fixing $\mathcal{D}(1)$ and $\mathcal{D}(2)$, analysis of the social welfare as a function of α and β reveals that it is maximized at $\alpha = \mathcal{D}(1), \beta = \mathcal{D}(2)$. This shows that in any case, no ranking function outperforms R_{PRP}. □

D. PROOF OF THEOREM 5

PROOF. Although this may sound intuitive, adding an action to players' action sets may actually decrease the social-welfare. We prove that whenever $c < 1/2$, the multi-topic document will never be a part of the support of the players. (The $c = 0.5$ case is technical and omitted due to lack of space.) This is done by showing it is *dominated*, i.e., it is never a part of the players' best-response. Let $s_2 \in \Delta([m])$ be a mixed strategy for the second player and let $j, \ell \in [m]$ be two distinct topics. The first publisher payoffs for writing on the topics are:

$$U_1(j, s_2) = \mathcal{D}(j) \cdot (1 - s_2(j)/2)$$
$$U_1(\ell, s_2) = \mathcal{D}(\ell) \cdot (1 - s_2(\ell)/2)$$

Without loss of generality, we assume $U_1(\ell, s_2) \geq U_1(j, s_2)$. The utility of writing the multi topic document is

$$U_1(\{j,\ell\}) = c \cdot (\mathcal{D}(j) \cdot (1 - s_2(j)/2) + \mathcal{D}(\ell) \cdot (1 - s_2(\ell)/2))$$
$$= c \cdot (U_1(j) + U_1(\ell)) < (U_1(j) + U_1(\ell))/2 \leq U_1(\ell).$$

This means that a publisher could always profit more from writing a single-topic document. Thus, publishers will never write on multiple topics in an equilibrium. □

E. PROOF OF EXAMPLE 1

PROOF. The payoff matrix of the $MTPG_{PRP}$ game is

$(1/3, 1/3)$	$(2/3, c/3)$	$(2/3, 1/3)$
$(c/3, 2/3)$	$(c/2, c/2)$	$(2c/3, 1/3)$
$(1/3, 2/3)$	$(1/3, 2c/3)$	$(1/6, 1/6)$

,

where for each player the first action is associated with writing a document on topic 1, the second is writing a multiple topic document, and the last action is writing on the second topic. For any value of c, there exists an equilibrium where both players write on topic 1. This equilibrium has a social welfare of 2/3, while an optimum of 1 is reached if both players write on different topics. This means that for any value of c, the price of anarchy of R_{PRP} is 1.5.

For $c > 1/2$, consider the ranking function R_{prob}, which is identical to R_{PRP}, except for the case where one publisher wrote about topic 1, the other about a multiple topics document, and a query on the first topic was posted. In this case, R_{prob} ranks the multi-topic document first with probability $1 - \frac{3c}{2(1+c)}$. The probabilistic ranking changes the utility of the players which are aiming to maximize their expected payoff, and the resulting payoff matrix is

$(1/3, 1/3)$	$\left(\frac{c}{1+c}, \frac{c}{1+c}\right)$	$(2/3, 1/3)$
$\left(\frac{c}{1+c}, \frac{c}{1+c}\right)$	$(c/2, c/2)$	$(2c/3, 1/3)$
$(1/3, 2/3)$	$(1/3, 2c/3)$	$(1/6, 1/6)$

.

This game has an optimal pure-strategies equilibria when players write on different documents, and a single mixed strategies equilibrium. In this mixed equilibrium, each player creates a document about the first topic with probability $\frac{3(c-1)c}{3c^2-7c+2}$, and a multi topic document otherwise. The resulting social welfare in this case is $\frac{2c(c^2-4c+1)}{(c+1)(3c^2-7c+2)}$, which approaches 1 when c is almost 1. □

F. PROOF OF THEOREM 7

PROOF. In $MTPG_{PRP} = (2, 2, c, \mathcal{D}; R_{PRP})$, there exists a dominant-strategy equilibrium in which both players write on the first topic. Playing it, the publishers earn $U(1) = p/2$, while the utility of writing about multiple topics is $U(\{1, 2\}) = c \cdot (1-p) \leq 1-p \leq p/2$, and the utility of writing about the second topic is $U(2) = 1 - p \leq p/2$. This means that the social welfare achieved by R_{PRP} is p. Next, we define R_{prob} as follows:

- Consider a query on the first topic, a document written on the topic, and a multi-topic document. In this case, R_{prob} will rank the multi-topic document first with probability $r \triangleq 1 - \frac{c}{p\cdot(1+c)}$. Notice that since $c \in (1/2, 1]$ and $p \in [2/3, \frac{2c}{c+1})$, we have $r \in [\frac{1}{4}, \frac{1}{2}]$.
- On all other cases, R_{prob} behaves just like R_{PRP}.

The introduction of probabilistic ranking in R_{prob} changes the game payoffs when one player writes on topic 1, while the other on multiple topics. In this case, the utility of writing on topic 1 drops to $U(1, \{1, 2\}) = p(1-r) = \frac{c}{1+c}$. Similarly, the utility of writing multiple topics rises to $U(\{1, 2\}, 1) = c(1 - p + p \cdot r) = \frac{c}{1+c}$. The game payoff matrix is therefore:

$(p/2, p/2)$	$\left(\frac{c}{1+c}, \frac{c}{1+c}\right)$	$(p, 1-p)$
$\left(\frac{c}{1+c}, \frac{c}{1+c}\right)$	$(c/2, c/2)$	$(c \cdot p, 1-p)$
$(1-p, p)$	$(1-p, c \cdot p)$	$\left(\frac{1-p}{2}, \frac{1-p}{2}\right)$

First, we notice that the scenario in which both players write on the first topic is no longer an equilibrium, as players are better off deviating to the multi-topic document, earning $\frac{c}{1+c} > p/2$. Next, the strategy of writing on topic 1 still dominates creating a document about the second topic. Finally, we observe that the pure-strategies equilibria ((1, 2) or (2, 1)) gives a social welfare of $\frac{2c}{1+c} > p$. The only mixed strategies equilibrium in this game is reached when each player writes about topic 1 with probability $x = \frac{c\cdot(c-1)}{c^2+c\cdot(p-3)+p}$ and otherwise writes a multi-topic document. This gives a social welfare of

$$P = 2U(2) = 2x \cdot \frac{c}{1+c} + (1-x) \cdot c$$
$$= \frac{c}{1+c} \cdot \frac{c^2p + 2cp - 4c + p}{c^2 + c \cdot (p-3) + p}$$

which, on the domain of $c \in (1/2, 1]$ and $p \in [2/3, \frac{2c}{c+1})$ is strictly larger than p. □

G. PROOF OF THEOREM 8

PROOF. We prove the theorem by showing that any equilibrium is associated with a matching in the bipartite graph which has the publishers on one partite and the topics on the other. Formally, let $G = ([n] \cup [m], [n] \times [m])$ be a full bipartite graph. Let $w : [n] \times [m] \rightarrow [0, 1]$ be a weight function such that $w(i, j) = Q_{i,j}$. We prove that an equilibrium for the game corresponds to a matching that could be generated by the greedy algorithm for graph matching. The algorithm iteratively expands the matching by simply adding the maximum-weight edge connecting unmatched vertices to the matching. For simplicity, the proof assumes $n \leq m$, although this is not required, as otherwise the utility of $n - m$ players is 0, and they may be removed to obtain a valid matching. We complete the proof by induction over the number of players in the game.

- **Basis:** in a single publisher game, the only equilibrium is when the player writes a document on the topic j maximizing $w(1, j)$, and so does the greedy matching.
- **Induction Step:** Let $i, j \in \text{argmax}\{Q_{i,j}\}$ be a top-quality publisher-topic pair. Since no other player has the same writing quality on topic j, player i writes about j, knowing he will surely be ranked first. The remaining players then know that no gain would come from writing on topic j. Hence, we are left with a game with $n - 1$ players and $m - 1$ topics. In the graph, this pair represents a maximum weight edge, and therefore may be selected by the greedy algorithm. Using the induction hypothesis, we conclude that the remaining players will also follow the greedy matching algorithm.

We conclude the proof by using the well-known result stating that the greedy algorithm achieves a 2-approximation for the maximum weight matching of the graph. □

H. PROOF OF THEOREM 9

PROOF. Consider the following Quality Matrix

players/topics	1	2	...	m
1	1	$1 - \epsilon/2$		0
2	$1 - \epsilon/2$	0		0
⋮			⋱	
n	0	0		0

All zeros Quality Matrix expect for $(1, 1), (1, 2), (2, 1)$.

Writing a document on topic 1 is a dominating strategy for player 1. Once this is done, the social welfare of the game is 1, regardless of the choice of the remaining players. The price of anarchy is as required, since the optimum is reached when players 1 and 2 write about topics 2 and 1 respectively yielding a social welfare of $2 - \epsilon$. □

I. PROOF SKETCH OF THEOREM 10

PROOF. Since price of anarchy analysis concerns with the *ratio* between the optimal strategy and an equilibrium strategy, it stays the same if $\mathcal{Q}$ is multiplied by a constant. Therefore, we assume that $\mathcal{Q}_{1,1} = 1$, i.e. the first publisher writes about the first topic perfectly. Therefore, a general Quality Matrix can be described using three parameters, $a, b, c \in [0, 1]$:

players/topics	1	2
1	1	a
2	b	c

The Quality Matrix $\mathcal{Q}$.

Notice that this notation means that the optimal solution for price of anarchy analysis is $OPT = \max\{1 + c, a + b\}$. We define the ranking function R_{uni} as the function which gives equal probability to all documents written on the topic whose quality is at least 2/3 times the quality of the best document on the topic. The proof is done by an extensive case analysis over the possible values of a, b and c, and is omitted due to lack of space. □

J. PROOF OF THEOREM 11

PROOF. Intuitively, we show that $\mathcal{R}$ has to be at least 1.5 worse than the optimum in one of two cases. If it prefers a "$\frac{2}{3}$-quality" document over a perfect document with probability larger than third, then both players might write on the topic they are less qualified to write on, thereby incurring a 1.5 anarchy. Alternatively, if $\mathcal{R}$ does not rank the worse document with that probability, players might selfishly write on their topic of choice, and no one would satisfy the information need of the remaining topics.

We denote $p \triangleq \mathcal{R}(2/3, 1)(1)$, i.e. the probability $\mathcal{R}$ ranks a document with quality 2/3 higher than a document of quality 1. Next, we consider two cases:

- $\mathbf{p \geq 1/3}$. In this case, consider the Quality Matrix:

players/topics	1	2
1	1	2/3
2	2/3	1

This means that the payoff for the players can be represented as:

$(1 - p, 2p/3)$	$(1, 1)$
$(2/3, 2/3)$	$(2p/3, 1 - p)$

. Since $1/3 \leq p \leq 1$, there exists an equilibrium in which player 1 writes about topic 2 while player 2 writes about topic 1. In this equilibrium, the payoff for each player is 2/3, while the optimum is 2. This means that the price of anarchy of $\mathcal{R}$ is at least $\frac{OPT}{P} = \frac{2}{4/3} = 1.5$.

- $\mathbf{p<1/3}$. When $\mathcal{R}$ gives high probability of ranking first the second-best document, it will be far from optimal in the following situation. Consider the Quality Matrix:

players/topics	1	2
1	1	$1 - p$
2	2/3	0

The following payoff matrix is then:

$(1 - p, 2p/3)$	$(1, 0)$
$(1 - p, 2/3)$	$(1 - p, 0)^4$

.

Notice that both players writing on the first topic is an equilibrium for the game, while the optimum is reached at 2, 1, thus $\frac{OPT}{P} = \frac{1-p+2/3}{1-p+\frac{2p}{3}} = \frac{5-3p}{3-p} \geq_{(p<1/3)} 1.5$.

□

[4]This implicitly assumes that a positive-quality document will be ranked higher than one with quality 0. The theorem holds regardlessly, as the payoff of this scenario is irrelevant.