



**HAL**  
open science

# Revenue-Maximizing Rankings for Online Platforms with Quality-Sensitive Consumers

Pierre L'Ecuyer, Patrick Maillé, Nicolas Stier-Moses, Bruno Tuffin

► **To cite this version:**

Pierre L'Ecuyer, Patrick Maillé, Nicolas Stier-Moses, Bruno Tuffin. Revenue-Maximizing Rankings for Online Platforms with Quality-Sensitive Consumers. 2014. hal-00953790v1

**HAL Id: hal-00953790**

**<https://inria.hal.science/hal-00953790v1>**

Preprint submitted on 28 Feb 2014 (v1), last revised 5 Apr 2016 (v3)

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Revenue-Maximizing Rankings for Online Platforms with Quality-Sensitive Consumers

Pierre L'Ecuyer\*, Patrick Maillé†, Nicolás Stier-Moses‡, and Bruno Tuffin§

## Abstract

When a customer searches for a keyword at a classified ads website, at an online retailer, or at a search engine (SE), the platform has exponentially many choices in how to sort the output to the query. The two extremes are (a) to consider a ranking based on relevance only, which attracts more customers in the long run because of perceived quality, and (b) to consider a ranking based on the expected revenue to be generated by immediate conversions, which maximizes short-term revenue. Typically, these two objectives are not perfectly positively correlated and hence the main question is what middle ground between them should be chosen. We introduce stochastic models and propose effective solution methods that can be used to optimize the ranking considering long-term revenues. A key feature of our model is that customers are quality-sensitive and are attracted to the platform or driven away depending on the average relevance of the output. The proposed methods are of crucial importance in e-business and encompass: (i) classified ad websites which can favor paid ads by ranking them higher, (ii) online retailers which can rank products they sell according to buyers' interests and/or the margins these products have, (iii) SEs which can position the content that they serve higher in the output page than third-party content to keep users in their platforms for longer and earn more. This goes in detriment of just offering rankings based on relevance only and is directly linked to the current *search neutrality debate*.

## 1 Introduction

The Internet occupies an increasingly important position in our daily lives. Electronic commerce, in particular, has enabled marketplaces in which participants can buy, sell, or rent a huge variety of objects and services in a very convenient way. Because the current Internet is a complex ecosystem of companies, there are various business models that have proved profitable. Among them, we will specially focus on three specific classes: online retailers such as Amazon.com that act as intermediaries between producers and consumers; classified ad websites such as eBay that allow sellers or service-providers, and buyers or service-consumers, respectively, to meet and conduct transactions; and search engines (SE) such as Google that allow users to find content of their interest on the web, and use these transactions as a chance to sell advertisement. To be profitable,

---

\*Département d'Informatique et de Recherche Opérationnelle, Université de Montréal, C.P. 6128, Succ. Centre-Ville, Montréal, H3C 3J7, Canada; and Inria Rennes – Bretagne Atlantique. E-mail: [lecuyer@iro.umontreal.ca](mailto:lecuyer@iro.umontreal.ca)

†Telecom Bretagne, 2 rue de la Châtaigneraie CS 17607, 35576 Cesson Sévigné Cedex, France. E-mail: [patrick.maille@telecom-bretagne.eu](mailto:patrick.maille@telecom-bretagne.eu)

‡Universidad Torcuato Di Tella Business School, Saenz Valiente 1010, Buenos Aires, Argentina; and CONICET Argentina. E-mail: [nstier@utdt.edu](mailto:nstier@utdt.edu)

§Inria Rennes – Bretagne Atlantique, Campus Universitaire de Beaulieu, 35042 Rennes Cedex, France. E-mail: [btuffin@irisa.fr](mailto:btuffin@irisa.fr)

those marketplaces typically rely on one or more of the following revenue streams. In some cases, they charge a commission equal to a percentage of the agreed price-tag (e.g., eBay or Airbnb). Some marketplaces provide a basic service for free but charge sellers to display their classified ads in premium locations or for additional time (e.g., leboncoin.fr in France, or Mercado Libre in Latin America). In addition, they offer additional services such as insurance or delivery for a fee. Finally, another common revenue source comes from third-party advertisers that display text or banners within the pages of the marketplace in exchange for payment.

The common feature in all those platforms is that when a user connects to them and inputs a query or category, the site provides a list of relevant items that match what the user wants. To provide value to users, it is crucial to present the relevant items in the platform in the correct order so the user can find the most appropriate ones. Indeed, by presenting certain items first, the site can boost users' interest by increasing relevance. For example, eBay provides relevance-based ranking, among other possibilities such as time until the end of the auction, distance, price, etc. The details of how to assign a relevance value to a query vary depending on the intrinsic details of the platform. For example, eBay may use the rating of the seller, Amazon may use the number of conversions for a product, and Google may use the click-through-rate as inputs.

Alternatively, the platform can take a more myopic approach and increase short-term revenue by favoring highly-profitable items. The key issue we shall address is how to achieve a balance between both goals when customers are quality-sensitive and their likelihood of visiting the platform is a function of the perceived relevance. Our results provide optimal ranking policies with respect to long-term revenue maximization. Also, we compare the optimal policy to other possible rankings—such as those based on relevance only or those based on short-term revenue only—in terms of expected revenue for the platform, expected revenue for the various content providers, and consumer welfare (captured by the expected quality).

Focusing now on SEs, companies in this segment have to rank the documents available on the Web that match the keyword entered by the user. A set of documents is selected and a relevance weight is assigned to each one using a link analysis algorithm [3, 4]. In addition to this list (usually referred to as *organic search*), the SE displays paid ads (commonly found at the top and at the right of the organic, unpaid, results). The most common arrangement is that advertisers pay the SE whenever users click on their ads. The payment amount is automatically decided using a bidding process between the SE and all advertisers interested in that keyword.

In the last couple of years, some SEs have been under scrutiny by individuals and organizations that oversee the Internet as well as by regulators in various countries because some believe that the organic search ranking is not only done with respect to an objective measure of relevance but that some revenue-making ingredients also play a role [10]. For example, it has been said that Google may favor YouTube content because of the extra revenue it generates [5, 26]. This debate is usually referred to as *search (non)neutrality* and it has ignited public interest [10, 15]. By definition a neutral SE should only use relevance to construct its rankings. This would allow new entrants that perform well (i.e., that are commonly clicked) to be listed near the top of the list of organic search results. The risk of a non-neutral ranking is that it may slow down innovation by favoring the incumbents that are known to generate profits, thereby preventing new applications/content from being shown, and hence to become known and successful. The term *search neutrality* has been first used in 2009 by Adam Raff, co-founder of Foundem (a price-comparison company, see <http://www.foundem.co.uk>), after a dispute with Google who penalized this company when showing search results in organic search and ads in Google's Adwords advertisement platform. It has been

reported that search bias was observed in experiments; for a general discussion on this issue, see [12], [28], and [20, Chapter 7].

Such discussions are natural extensions of the current ones on *network neutrality*, which is the topic of an extensive debate worldwide. The debate is about whether Internet Service Providers should treat all bits in the same way or some traffic could be granted a priority status over the rest, and whether payment is an acceptable mechanism to allocate priorities. For an introduction to the topic, we refer the reader to [17, 19, 22]; for a more technical approach based on mathematical modeling and game theory, see [1, 7, 9, 13, 21]. Although both debates are deeply connected, the search neutrality one has not yet attracted much attention in the literature. It has been mentioned in [10], and in Inria’s response to the ARCEP (the French regulator) national consultation on network neutrality [15]. A preliminary model related to that presented here has been developed in [8]. Finally, there are other initiatives that model the economics of SEs, but they do not explicitly investigate the impact of a non-neutral ranking [24].

Motivated by this debate, in this article we consider ranking policies and we study their impact on the marketplace or SE revenue, as well as on social welfare. Two usually opposing effects need to be considered:

- If the ranking designer considers a short time-horizon, it may prefer to place items that yield the largest immediate expected gain on top. For example, Google normally places YouTube content above other third-party video content. When sorting on relevance, a retailer may choose to display more expensive items before others because those items provide higher revenues.
- From the customer point-of-view, an optimal ranking should be based on relevance. Indeed, customer welfare is maximized by sorting items by decreasing order of likelihood to please customers. Those likelihoods are estimated based on quality, customer’s interest, and price if any. This ranking is also compatible with the view of a ranking designer that considers a long time-horizon since maximizing perceived quality attracts more visitors to the site, generating more revenue over the long run. To illustrate, Google’s search engine algorithm has been perceived since its early times as very quick and effective. This perception of quality won them an almost monopolistic status worldwide. For example in late 2013, it had about 67% marketshare in the search market, although it was much bigger in some countries (e.g., 90% in France). For more statistics, see comScore’s website, and [25].

The main goal of this paper is to develop a modeling framework that permits us to design tractable algorithms for computing optimal ranking policies for the platform, assuming that customers are quality-sensitive and may defect to competing platforms if they do not find what they are looking for. Surprisingly, as far as we know, we are the first to provide an economic analysis of ranking policies, and to show how to design optimal policies from the perspective of the platform. The tools we develop can prove useful to websites that want to fine-tune ranking policies to achieve long-term profitability. Since SEs play an important role in our connected society by allowing end-users to access content and applications without necessarily knowing of them, such revenue-maximizing strategies are directly linked to the search neutrality debate. Hence, the framework we introduce can also be of high interest to regulators who study the impact that search neutrality has on users and on overall social welfare. This may allow regulators to determine if intervention is warranted, and to study the consequences of doing so. In particular, our study may prove useful to

provide arguments for or against non-neutral SEs. We will investigate the impact that considering different rankings has on the various participants.

Ranking items or links based on relevance in the context of SEs has been the topic of extensive academic research (see, e.g., [3, 4, 16, 27] and the references therein for details). The main purpose of these studies is defining how to best compute a relevance index that is appropriate to rank pages. Instead, our focus is the procedure itself that is used to rank items; henceforth, in this article we shall assume that relevance is part of the input data.

In our model, requests arrive to the site over time at a rate that depends on the average relevance of displayed links; the more relevant the expected results, the larger the number of visits. The expected revenue is the rate of visits to each page, multiplied by the expected revenue per visit for that page, summed over all possible pages. These quantities depend on the *ranking policy*, defined as a rule that assigns a permutation of matching pages to each possible keyword. A ranking based only on immediate revenue is generally suboptimal; it must also take into account the impact on relevance because that affects the rate of visits since customers are quality-sensitive. For the purpose of this study, we consider that the distributions of relevance and expected revenues for each page are known in advance, so we consider them as inputs. (These distributions can be estimated empirically using data available in the SEs' and marketplaces' servers. Although exploring this data is an interesting direction of research, we leave this to follow-up work.) Our main contribution is the characterization of the optimal policy for the ranking problem, which turns out to be a complex task, in part because the objective function is not additive, and classical dynamic programming tools cannot be used directly. We propose an algorithm that exploits the characterization of the optimal policy that allows the SE to assign a scalar number to each matching item to then find the optimal ranking by simple sorting.

The rest of the article is organized as follows. Section 2 presents our modeling framework while Section 3 explains how the ranking problem can be simplified so one does not need to consider the exponentially-many possible orderings. Using the conditions presented there, we show how it suffices for a SE to sort the pages with respect to a scalar number, coming from a linear combination of relevance and revenue. Having characterized optimal rankings, in Section 4 we present an algorithm that computes the correct linear combination of relevance and revenue, which enables the SE to execute the sorting procedure. Section 5 discusses the impact that arises from an SE that takes a middle ground between offering a search-neutral output and a myopic one considering only short-term revenue. Finally, we offer conclusions in Section 6.

## 2 Model Formulation

In this section we provide the definition of the model we consider. For the presentation, we use the context of a SE that receives keywords and generates a list of organic results using links to relevant and/or profitable web pages. The model could be easily adapted to be used by other marketplaces such as electronic retailers and classified-ad websites.

For each arriving request (i.e., a query sent to the SE by a user), different content providers (CPs) host pages that are relevant. Out of a universe of  $m_0$  pages available online, we denote by  $M \leq m_0$  the number of pages that match the arriving request. Each page  $i = 1, \dots, M$  has a relevance value  $R_i \in [0, 1]$ , and an expected revenue per click  $G_i \in [0, K]$  for the CP (here,  $K$  is a positive constant) of which the SE receives a fraction  $\alpha_i \in [0, 1]$ . Consequently, the SE's expected revenue per click from page  $i$  is  $\alpha_i G_i$ . The SE might sometimes also be the CP for a subset of the

pages matching the request; in those cases  $\alpha_i = 1$  because it receives all the revenue. Putting this all together, the instance of the ranking problem corresponding to a given request is encoded by a vector  $Y = (M, R_1, G_1, \alpha_1, \dots, R_M, G_M, \alpha_M)$  that we assume to belong to a universe of admissible requests. After getting the request, the SE must select a permutation  $\pi = (\pi(1), \dots, \pi(M))$  of the  $M$  pages and use it to display links to those pages in order. A stationary *ranking policy*  $\mu$  is a function that assigns a permutation  $\pi = \mu(Y)$  to each possible realization of  $Y$ . Except when otherwise indicated, we shall only consider *deterministic* stationary policies, as opposed to *randomized* ones, which map each  $Y$  to a probability distribution over the set of permutations of  $M$  elements.

The *click-through-rate* (CTR) of a link that points to a page is defined as the probability that the user clicks on that link [14, Chapter 8]. This probability depends on the relevance of the content but also on the position number where the link is displayed. We assume that the CTR of the link to page  $i$  placed at position  $\pi(i)$  can be expressed as a combination of a position effect and a relevance effect. That is, CTR is given by

$$\text{CTR}(i) = \theta_{\pi(i)} \psi(R_i),$$

where  $1 \geq \theta_1 \geq \theta_2 \geq \dots \geq \theta_m > 0$  is a non-increasing sequence of fixed positive constants that describe the importance of each position in the ranking. The non-decreasing function  $\psi : [0, 1] \rightarrow [0, 1]$  maps the relevance to the (position-independent) probability of the page. The assumption that the CTR is separable is pervasive in the e-Commerce literature [18]. We will rely on it to derive simple optimality conditions. According to this assumption, to increase the CTR, we can either choose a more relevant page or we can choose a position closer to the top of the list.

Fixing a request  $Y$  and a permutation  $\pi$ , we now define the various objective functions we shall consider. The *local relevance* captures the attractiveness of the ordering from the consumer's perspective. It is computed by

$$r(\pi, Y) := \sum_{i=1}^M \text{CTR}(i) R_i = \sum_{i=1}^M \theta_{\pi(i)} \psi(R_i) R_i = \sum_{i=1}^M \theta_{\pi(i)} \tilde{R}_i, \quad (1)$$

where  $\tilde{R}_i := \psi(R_i) R_i$ . The expected total revenue arising from the request equals

$$g_0(\pi, Y) := \sum_{i=1}^M \text{CTR}(i) G_i = \sum_{i=1}^M \theta_{\pi(i)} \psi(R_i) G_i, \quad (2)$$

out of which the SE receives

$$g(\pi, Y) := \sum_{i=1}^M \text{CTR}(i) \alpha_i G_i = \sum_{i=1}^M \theta_{\pi(i)} \psi(R_i) \alpha_i G_i = \sum_{i=1}^M \theta_{\pi(i)} \tilde{G}_i, \quad (3)$$

where  $\tilde{G}_i := \alpha_i \psi(R_i) G_i$ .

To obtain an optimal ranking policy, we must consider that since customers are quality-sensitive, the choice of policy  $\mu(\cdot)$  influences the future arrivals of customers. This has deep implications because a myopic policy for the SE (i.e., choosing  $\mu(Y) \in \arg \max_{\pi} g(\pi, Y)$  for each  $Y$ ) does not suffice to achieve optimality. To capture the dependence on future end-users that arrive to the SE, we consider the multivariate distribution of the input requests  $Y$ . Each request is then interpreted as a realization of  $Y$  according to that distribution.

We estimate the long-term value induced by a stationary ranking policy  $\mu$  by taking expectations of the objectives presented earlier with respect to the distribution of input requests. Therefore, the expected relevance per request is

$$r := r(\mu) = \mathbb{E}[r(\mu(Y), Y)], \quad (4)$$

the expected total revenue per request is

$$g_0 := g_0(\mu) = \mathbb{E}[g_0(\mu(Y), Y)], \quad (5)$$

and the expected SE revenue per request is

$$g := g(\mu) = \mathbb{E}[g(\mu(Y), Y)]. \quad (6)$$

In the three previous definitions, the expectation is taken with respect to the random variable  $Y$ .

As discussed in the introduction, a non-myopic SE would be interested in the expected long-run revenue. This must depend on both the expected relevance per request  $r$  and on the expected SE revenue per request  $g$ . We capture the two dependencies through the general function

$$U_{\text{SE}} = \varphi(r, g), \quad (7)$$

where  $\varphi$  is an increasing function of  $r$  and  $g$  with bounded second derivatives over  $[0, 1] \times [0, K]$ . An *optimal policy* from the perspective of the SE is a stationary ranking policy  $\mu$  that maximizes  $U_{\text{SE}}$ .

We are going to pay special attention to the class of ranking policies that sort the  $M$  pages by decreasing order of their value of  $\tilde{R}_i + \rho\tilde{G}_i$ , for a given constant  $\rho > 0$ . We refer to such policy as a *linear ordering (LO) policy with ratio  $\rho$*  (or LO- $\rho$  policy, for short). In fact, if  $\theta_k = \theta_{k+1}$ , the ordering at positions  $k$  and  $k + 1$  does not matter, and we still say that we have an LO- $\rho$  policy regardless of the order at these positions. When  $\rho = 0$ , the ordering is based only on  $\tilde{R}_i$ , whereas in the limit as  $\rho \rightarrow \infty$ , the ordering is based only on  $\tilde{G}_i$ . We show below that under mild conditions on the request distribution, an optimal policy  $\mu^*$  coincides with an LO- $\rho$  policy for a specific value of  $\rho$  that we will characterize. We also highlight that specifying  $\rho$  is not enough to uniquely characterize an optimal policy in the case when several  $\tilde{R}_i + \rho\tilde{G}_i$  may be equal with positive probability.

The objective function (7) is very general, and the assumptions written after its definition are enough to develop our theoretical analysis. In practice  $\varphi(r, g)$  usually takes the form of an (average) arrival rate of requests multiplied by an expected revenue per request. To develop intuition, the examples we provide below have more structure, as we now describe. We assume that search requests arrive according to a (Poisson) process of (constant) rate  $\lambda(r)$ , where  $\lambda : [0, 1] \rightarrow [0, \infty)$  is an increasing, positive, smooth (continuously differentiable), and bounded function. Its argument  $r$  is the average relevance corresponding to the policy in use, as defined in (4). Each time the SE receives a search request, it gets a revenue  $\beta$  in expectation from the third-party advertisement displayed in the page. Hence, the expected SE advertisement revenue per time unit is  $\beta\lambda(r)$ , which depends on the ranking policy only via  $r$ . On top of this, and as discussed earlier, we assume that the SE receives a proportion of the CP revenue, totalling  $\lambda(r)g$ . Putting it all together, the total expected SE revenue per unit time in our examples is

$$U_{\text{SE}} = \lambda(r)(\beta + g). \quad (8)$$

This expression is increasing in  $r$  and in  $g$ . When  $\alpha_i = 0$  for all  $i$ ,  $g = 0$  and the SE's best interest is being neutral to just maximize  $r$ , i.e., to rank based on relevance. Otherwise, if  $g > 0$ , the SE may be interested in selecting permutations that increase  $g$  even if this decreases  $r$  a bit. The larger the  $\alpha_i$ 's, the stronger the incentive of the SE to consider non-neutral ranking policies. In the next sections, we characterize optimal strategies and develop algorithms to compute or approximate them.

### 3 Optimality Conditions for Ranking Policies

In this section we derive optimality conditions on the permutation  $\pi = \mu(Y)$  associated with any given request vector  $Y$ . Later, these conditions will be used to develop computational algorithms that can provide a ranking for each  $Y$ . We first develop *approximate* necessary optimality conditions under the assumption that  $Y$  has a discrete distribution. Then we show that these conditions must hold in the limit if we assume that each  $Y$  has a negligible probability, that is, if  $Y$  is a continuous random variable with a density. Under further assumptions, these necessary conditions determine the optimal policy uniquely, up to a set of realizations  $Y$  of measure 0. This provides simple, approximate optimality conditions for the situation where the current request  $Y$  has a small probability, small enough that changing  $\pi$  in the local relevance  $r(\pi, Y)$  can only bring a small change to  $r$  and  $g$ .

#### 3.1 Necessary Optimality Conditions Under a Discrete Distribution for $Y$

Let us first suppose that  $Y$  has a discrete distribution  $p(y) = \mathbb{P}[Y = y]$ . Assume that  $\mu$  is an optimal policy, with  $r$  and  $g$  the associated objectives. Since  $\mu$  is optimal, for any  $y$ , permuting two successive elements in  $\pi = \mu(y)$ , say at positions  $k$  and  $k + 1$ , must not increase the expected long-term revenue. Let  $\delta = \pi^{-1}$ , the inverse permutation to  $\pi$ . Then the numbers of the pages at positions  $k$  and  $k + 1$  are  $\delta(k)$  and  $\delta(k + 1)$ . Let  $\Delta\theta = \theta_k - \theta_{k+1} \geq 0$ . Switching the two pages at positions  $k$  and  $k + 1$  will permute the vectors  $(\tilde{R}_{\delta(k)}, \tilde{G}_{\delta(k)})$  and  $(\tilde{R}_{\delta(k+1)}, \tilde{G}_{\delta(k+1)})$  in (1) and (3). The changes on  $r$  and  $g$  resulting from this switch would be

$$\Delta r = (\tilde{R}_{\delta(k+1)} - \tilde{R}_{\delta(k)})\Delta\theta p(y)$$

and

$$\Delta g = (\tilde{G}_{\delta(k+1)} - \tilde{G}_{\delta(k)})\Delta\theta p(y).$$

The corresponding change in  $U_{SE}$  is

$$\Delta U_{SE} = \varphi(r + \Delta r, g + \Delta g) - \varphi(r, g) = \varphi_r(r, g)\Delta r + \varphi_g(r, g)\Delta g + \mathcal{O}((|\Delta r| + |\Delta g|)^2),$$

where  $\varphi_r$  and  $\varphi_g$  are the partial derivatives of  $\varphi$  with respect to  $r$  and  $g$ , respectively. The optimality of  $\pi$  (or equivalently of  $\delta$ ) implies that this change on  $U_{SE}$  cannot be positive, so we must have

$$\varphi_r(r, g)\Delta r + \varphi_g(r, g)\Delta g \leq \mathcal{O}((|\Delta r| + |\Delta g|)^2),$$

yielding, whenever  $\Delta\theta p(y) > 0$ ,

$$\varphi_r(r, g)(\tilde{R}_{\delta(k+1)} - \tilde{R}_{\delta(k)}) + \varphi_g(r, g)(\tilde{G}_{\delta(k+1)} - \tilde{G}_{\delta(k)}) \leq \mathcal{O}(\Delta\theta p(y))$$



since  $\tilde{R}$  and  $\tilde{G}$  are bounded. This can be rewritten as

$$\varphi_r(r, g)\tilde{R}_{\delta(k+1)} + \varphi_g(r, g)\tilde{G}_{\delta(k+1)} \leq \varphi_r(r, g)\tilde{R}_{\delta(k)} + \varphi_g(r, g)\tilde{G}_{\delta(k)} + \mathcal{O}(\Delta\theta p(y)), \quad (9)$$

which must hold for all  $y$  and all  $k$  for which  $\Delta\theta p(y) > 0$ . If  $\Delta\theta p(y) = 0$ , there is no change on  $\Delta r$  or  $\Delta g$ , so the order at positions  $k$  and  $k + 1$  does not matter.

For every pair  $(r, g)$ , we set

$$h(r, g) := \frac{\varphi_g(r, g)}{\varphi_r(r, g)}, \quad (10)$$

for which we assume that  $\varphi_r(r, g) > 0$ . Using this notation, if  $p(y) \ll 1$  and we decide to neglect the  $\mathcal{O}(\Delta\theta p(y))$  term in (9), we obtain the following (approximate) *necessary* optimality conditions: When  $\theta_k > \theta_{k+1}$ , we must have

$$\tilde{R}_{\delta(k+1)} + h(r, g)\tilde{G}_{\delta(k+1)} \leq \tilde{R}_{\delta(k)} + h(r, g)\tilde{G}_{\delta(k)}. \quad (11)$$

These necessary conditions tell us that if  $\mu$  is an optimal policy and if the  $\mathcal{O}(\Delta\theta p(y))$  terms can be neglected, then  $\mu$  must be an LO- $\rho$  policy with ratio  $\rho = h(r, g)$ .

For the special case of the running examples introduced in the previous section, we have  $\varphi(r, g) = \lambda(r)(\beta + g)$ , which implies that  $\varphi_r(r, g) = \lambda'(r)(\beta + g)$ ,  $\varphi_g(r, g) = \lambda(r)$ , and

$$h(r, g) = \frac{\lambda(r)}{\lambda'(r)(\beta + g)}. \quad (12)$$

The conditions in (11) suggest that in a search for a (near-)optimal policy, we may restrict ourselves to LO- $\rho$  policies and try to optimize the value of  $\rho$ . This may appear simple at first sight, but there are potential difficulties with this plan. First, the  $\mathcal{O}(\Delta\theta p(y))$  term may be non-negligible, when  $p(y)$  is not very small. Second, finding the optimal  $\rho$  is not necessarily obvious or easy. Third, fixing  $\rho$  does not necessarily determine a unique policy, because there might be equalities in (11) and then the selected order might still matter. Fourth, when such equalities happen, there are situations where the optimal policy must be randomized (e.g., select one order with some probability  $p$  and the other with probability  $1 - p$ ; see below). Fifth, the conditions (11) are necessary for an optimal policy, but perhaps not sufficient. The following example illustrates those difficulties.

**Example 1.** We consider an instance with two pages and always a unique request. The input data consists of  $Y = (M, R_1, G_1, \alpha_1, R_2, G_2, \alpha_2) = (2, 1, 0, 0, 1/5, 8, 1/4)$  with probability 1,  $\psi(R) = 1$  for all  $R$ ,  $\lambda(r) = r$ ,  $\beta = 1$ , and  $(\theta_1, \theta_2) = (1, 1/2)$ . Replacing in the objective, we have  $\varphi(r, g) = r(1 + g)$ .

At each request, we must select a ranking, either (1, 2) or (2, 1). Suppose that instead of always selecting the same ranking for all requests, we adopt a policy that selects the ranking (1,2) with probability  $p$  and (2,1) with probability  $1 - p$ . We want to find the optimal value of  $p \in [0, 1]$ . For this randomized policy, we compute

$$\begin{aligned} r &= p(\theta_1 R_1 + \theta_2 R_2) + (1 - p)(\theta_1 R_2 + \theta_2 R_1) = (7 + 4p)/10, \\ g &= p(\theta_1 \alpha_1 G_1 + \theta_2 \alpha_2 G_2) + (1 - p)(\theta_1 \alpha_2 G_2 + \theta_2 \alpha_1 G_1) = 2 - p, \\ \varphi(r, g) &= r(1 + g) = (7 + 4p)(3 - p)/10 = (21 + 5p - 4p^2)/10. \end{aligned}$$

The objective function is quadratic and it attains its maximum at  $p^* = 5/8$ . Evaluating,  $r = 19/20$ ,  $g = 11/8$ , and  $\varphi(r, g) = 361/160$ . Note that by taking  $p = 0$  we get  $21/10 = 336/160$  and by taking

$p = 1$  we get  $22/10 = 352/160$ . This clearly shows that randomized ranking policies can perform better than deterministic ones.

Here we have  $h(r, g) = r/(1 + g) = (7 + 4p)/(10(3 - p))$ . With the optimal  $p^* = 5/8$ , this expression evaluates to  $h(r, g) = 2/5$ . If we consider the LO- $\rho$  rule with  $\rho = \rho^* = 2/5$ , we have  $\tilde{R}_1 + \rho^* \tilde{G}_1 = \tilde{R}_2 + \rho^* \tilde{G}_2 = 1$ . So with the  $\rho^*$  that corresponds to the optimal randomized policy, the ordering conditions are always satisfied, regardless of the order, because the two linear expressions are equal. On the other hand, this  $\rho^*$  is not sufficient to determine the optimal policy! Any policy satisfies the ordering conditions with  $\rho = \rho^*$ , but is not necessarily optimal. Moreover, as mentioned earlier, none of the two possible deterministic policies is optimal.

If we adopt an LO- $\rho$  policy with  $\rho \neq \rho^*$ , then the choice of  $\rho$  defines the policy uniquely, but this policy is not optimal either. Indeed, if  $\rho < 2/5$ , then we always take the order (1,2), which gives  $r = 11/10$ ,  $g = 1$ , and  $h(r, g) = 11/20 > 2/5$ , so the LO rule with  $\rho = h(r, g)$  tells us to always select the order (2,1). Reciprocally, if we always select the order (2,1), we obtain  $h(r, g) < 2/5$  and the LO rule with  $\rho = h(r, g)$  always tells us to select the order (1,2). Thus, we cannot guarantee the existence of optimal rankings with deterministic policies.  $\square$

Perhaps one could argue that the problem of this example comes from the fact that the instance is deterministic so  $p(Y) = 1$ . In fact, this is not the case. It is always possible to construct request densities that assign small probabilities  $p(y)$  to all request realizations  $y$ . This can be done by splitting artificially each possible realization of  $Y$  into an arbitrary large number of subrealizations, say  $\ell$ , each one having probability  $p(y)/\ell$ . Conceptually, one would achieve this by adding one artificial component to  $Y$  to obtain an extended vector  $Y'$  whose added component only identifies what subrealization we have. In this case, the ranking policy could output different rankings according to the subrealization  $Y'$  that was drawn. In other words, different permutations  $\pi$  can be selected for the same (original) realization  $Y$ . In the limit when  $\ell \rightarrow \infty$ , this mechanism effectively mimics a randomized ranking policy, where for any given realization  $Y$ , each permutation  $\pi$  is selected with a given probability. This randomized ranking policy is effectively specified as a deterministic policy in terms of the extended vector  $Y'$ , which in the limit when  $\ell \rightarrow \infty$  has a density (the artificial extra component has a continuous distribution).

Inspired by these observations, in the next subsection we study a framework in which  $Y$  is assumed to have a continuous distribution. This can be seen as an approximation when the set of possible request inputs is a huge set and each  $p(y)$  is increasingly small. We will provide conditions under which there is a non-randomized optimal policy, and then show how it can be computed.

### 3.2 Approximation by a Continuous Distribution for $Y$

In this section, we extend the discussion to include continuous distributions for the input requests. We let  $Y$  be a continuous random vector, with probability measure  $\nu$  over the class  $\Omega$  of Borel subsets of vectors  $Y = (M, R_1, G_1, \alpha_1, \dots, R_M, G_M, \alpha_M)$  where  $M \in \{1, \dots, m_0\}$  and  $(R_i, G_i, \alpha_i) \in [0, 1] \times [0, K) \times [0, 1]$  for each  $i$ . We assume that it has a (finite) density function  $f$ . That is, for each  $D \in \Omega$ ,  $\nu(D) = \int_D f(y) dy$ . Then, if  $\nu(D) > 0$ , we can always select  $\tilde{D} \subset D$ ,  $\tilde{D} \in \Omega$ , such that  $\nu(\tilde{D})$  is positive and arbitrary small.

Suppose that  $\mu$  is an optimal policy, with its corresponding  $r$  and  $g$ , and that we change  $\mu$  into  $\mu'$  by permuting the two successive elements at positions  $k$  and  $k + 1$  in  $\mu(y)$ , for all  $y \in D$ , for some fixed  $D \in \Omega$ . Note that  $\mu(y)$  might not be the same for all  $y \in D$ . Let  $\delta(k)(y)$  and  $\delta(k + 1)(y)$  be those page numbers at positions  $k$  and  $k + 1$  in  $\mu(y)$ , for each  $y$ . Switching those two pages

permutes the vectors  $(\tilde{R}_{\delta(k)(y)}, \tilde{G}_{\delta(k)(y)})$  and  $(\tilde{R}_{\delta(k+1)(y)}, \tilde{G}_{\delta(k+1)(y)})$ , for all  $y \in D$ . The changes on  $r$  and  $g$  coming from this switch are

$$\Delta r = \int_D (\tilde{R}_{\delta(k+1)(y)} - \tilde{R}_{\delta(k)(y)}) \Delta \theta f(y) dy$$

and

$$\Delta g = \int_D (\tilde{G}_{\delta(k+1)(y)} - \tilde{G}_{\delta(k)(y)}) \Delta \theta f(y) dy.$$

Since  $\Delta \theta \leq 1$ ,  $R_i \in [0, 1]$ ,  $\psi(R_i) \in [0, 1]$ ,  $G_i \in [0, K]$ , and  $\alpha_i \in [0, 1]$ , these changes satisfy  $|\Delta r| \leq \nu(D)$  and  $|\Delta g| \leq K\nu(D)$ . The corresponding change on  $\varphi$  is

$$\Delta U_{SE} = \varphi_r(r, g) \Delta r + \varphi_g(r, g) \Delta g + \mathcal{O}((|\Delta r| + |\Delta g|)^2). \quad (13)$$

We now rigourously define the linear ordering policies with ratio  $\rho$  that were mentioned in Section 2, and establish that optimal ranking policies belong to that category for a specific value of  $\rho$ .

**Definition 1.** A policy  $\mu$  is called an LO- $\rho$  policy if for almost all  $Y$  (with respect to the measure  $\nu$ ),  $\mu$  sorts the pages by decreasing order of  $\tilde{R}_i + \rho \tilde{G}_i$ , except perhaps at positions  $k$  and  $k+1$  where  $\theta_k = \theta_{k+1}$ , at which the order can be arbitrary. That is, whenever  $\theta_k > \theta_{k+1}$ , one must have

$$\tilde{R}_{\delta(k+1)(Y)} + \rho \tilde{G}_{\delta(k+1)(Y)} \leq \tilde{R}_{\delta(k)(Y)} + \rho \tilde{G}_{\delta(k)(Y)}. \quad (14)$$

**Proposition 1.** If the tuple  $(r, g)$  corresponds to an optimal policy, then this policy must be an LO- $\rho$  policy with  $\rho = h(r, g)$ .

*Proof.* The proof is by contradiction. Take a  $k$  such that  $\Delta \theta := \theta_k - \theta_{k+1} > 0$ , and suppose that there exists  $\epsilon > 0$  and  $D \in \Omega$  such that  $\nu(D) > 0$  and

$$\varphi_r(r, g) \tilde{R}_{\delta(k+1)(Y)} + \varphi_g(r, g) \tilde{G}_{\delta(k+1)(Y)} > \varphi_r(r, g) \tilde{R}_{\delta(k)(Y)} + \varphi_g(r, g) \tilde{G}_{\delta(k)(Y)} + \epsilon$$

for all  $Y \in D$ . Then,

$$\begin{aligned} \epsilon \nu(D) &\leq \int_D \left[ \varphi_r(r, g) \tilde{R}_{\delta(k+1)(y)} + \varphi_g(r, g) \tilde{G}_{\delta(k+1)(y)} - \varphi_r(r, g) \tilde{R}_{\delta(k)(y)} - \varphi_g(r, g) \tilde{G}_{\delta(k)(y)} \right] f(y) dy \\ &= (\varphi_r(r, g) \Delta r + \varphi_g(r, g) \Delta g) / \Delta \theta. \end{aligned}$$

The current policy being optimal, performing the permutation of positions  $k$  and  $k+1$  over  $D$  can only reduce the revenue, i.e.,  $\Delta U_{SE} \leq 0$ . From (13), for  $\nu(D)$  sufficiently small there exists a constant  $C$  such that

$$\varphi_r(r, g) \Delta r + \varphi_g(r, g) \Delta g \leq C((|\Delta r| + |\Delta g|)^2) \leq C(1 + K)^2 \nu^2(D).$$

Therefore,  $\nu(D) \geq \epsilon \Delta \theta / (C(1 + K)^2)$ . This also holds for  $D$  replaced by any  $\tilde{D} \subset D$  with  $\tilde{D} \in \Omega$ . By taking  $\tilde{D} \subset D$  small enough so that  $\nu(\tilde{D}) < \epsilon \Delta \theta / (C(1 + K)^2)$ , we obtain a contradiction. Therefore, (14) must hold with  $\rho = h(r, g)$  for all  $Y$  except perhaps over a set of measure zero.  $\square$

Proposition 1 tells us that any optimal policy must satisfy the LO- $\rho$  conditions for  $\rho = h(r, g)$ . But we need further assumptions to make sure that this specifies an optimal policy. In the rest of this section, we assume that the following condition holds.

**Assumption A.** For any  $\rho \geq 0$ , and any  $j > i > 0$ ,  $\mathbb{P}[M \geq j \text{ and } \tilde{R}_i + \rho \tilde{G}_i = \tilde{R}_j + \rho \tilde{G}_j] = 0$ .  $\square$

One example of a sufficient condition for this assumption to hold is that each pair  $(\tilde{R}_i, \tilde{R}_j)$  has a bivariate density with no probability mass concentrated on a set of Lebesgue measure zero in two dimensions (such as a line, for example).

Under Assumption A, for any fixed  $\rho \geq 0$ , there is an LO- $\rho$  policy  $\mu = \mu(\rho)$  that sorts almost any  $Y \in \Omega$  (i.e., there is a unique order with probability 1). This ranking policy has corresponding values of  $(r, g) = (r(\mu), g(\mu))$  and of  $h(r, g)$  that are uniquely defined. To refer to  $h(r, g)$  as a function of  $\rho$ , we write  $\tilde{h}(\rho)$ . From Proposition 1, if  $\mu$  is the optimal policy, then  $\rho$  must be a fixed point of  $\tilde{h}$ . Indeed, one must have  $\rho = h(r(\mu(\rho)), g(\mu(\rho))) = \tilde{h}(\rho)$ . The next proposition says that under certain conditions such a fixed point exists and is unique.

**Proposition 2.** (i) If  $h(r, g)$  is bounded over  $[0, 1] \times [0, K]$ , then the fixed-point equation

$$\tilde{h}(\rho) = \rho \tag{15}$$

has at least one solution in  $[0, \infty)$ .

(ii) If the derivative  $\tilde{h}'(\rho) < 1$  for all  $\rho > 0$ , then the solution is unique.

*Proof.* (i) If  $\tilde{h}(0) = 0$ , then  $\rho = 0$  is already a solution. Otherwise, we have  $\tilde{h}(0) > 0$  and  $\tilde{h}(\rho) \leq K'$  for all  $\rho \geq 0$  for some constant  $K'$ . In particular,  $\tilde{h}(K') \leq K'$ . Since  $\tilde{h}$  is a continuous function, there must be at least one point  $\rho \in [0, K']$  at which  $\tilde{h}(\rho) = \rho$ .

(ii) The slope of the function  $\tilde{h}(\rho)$  is always less than 1, so it cannot cross the line  $f(\rho) = \rho$  more than once.  $\square$

Going back to the special case of our running examples where  $\varphi(r, g) = \lambda(r)(\beta + g)$ , Proposition 2 becomes:

**Proposition 3.** Suppose  $\varphi(r, g) = \lambda(r)(\beta + g)$ .

(i) If  $\lambda(r)/\lambda'(r)$  is bounded for  $r \in [0, 1]$  and  $\beta + g(0) > 0$ , then (15) has at least one solution in  $[0, \infty)$ .

(ii) If  $\lambda(r)/\lambda'(r)$  is also non-decreasing in  $r$ , then the solution is unique.

*Proof.* (i) In this case, we have

$$\tilde{h}(\rho) = \frac{\lambda(r(\mu(\rho)))}{\lambda'(r(\mu(\rho)))(\beta + g(\mu(\rho)))}.$$

Note that  $g(\mu(\rho)) \geq g(\mu(0))$  for all  $\rho \geq 0$ . Therefore, the conditions in (i) imply that  $\tilde{h}(\rho)$  is bounded and we can apply Proposition 2 (i).

(ii) If  $\lambda(r)/\lambda'(r)$  is non-decreasing in  $r$ , then it is non-increasing in  $\rho$  since  $r(\mu(\rho))$  is non-increasing in  $\rho$ . Additionally, since we know that  $g(\mu(\rho))$  is non-decreasing in  $\rho$ , it follows that  $\tilde{h}(\rho)$  is non-increasing in  $\rho$ , so  $\tilde{h}'(\rho) \leq 0$  and we can apply Proposition 2 (ii).  $\square$

The condition that  $\lambda(r)/\lambda'(r)$  is non-decreasing, in (ii), is actually a bit stronger than what we need to satisfy the condition of Proposition 2 (ii). To illustrate when this condition is satisfied, take  $\lambda(r) = a_0 + b_0 \ln(c_0 + r)$  for some constants  $a_0 \geq 0$ ,  $b_0 > 0$ , and  $c_0 \geq 1$ . Then,  $\lambda'(r) = b_0/(c_0 + r)$ , and therefore  $\lambda(r)/\lambda'(r) = [a_0 + b_0 \ln(c_0 + r)](c_0 + r)/b_0$ , which is bounded and increasing in  $r \in [0, 1]$ . Other simple cases where the condition holds are the monomial forms  $\lambda(r) = a_0 r^{b_0}$  for any positive values  $a_0$  and  $b_0$ ; which includes the case  $\lambda(r) = r$  considered in several examples in this paper.

When  $\alpha = 0$ , we immediately find  $\rho^* = 0$  and Proposition 3 simplifies to the following intuitive result, which establishes that it is optimal for the SE to rank according to relevance. In this case, the SE has the incentive to conform to search neutrality.

**Corollary 1.** *If  $\varphi(r, g) = \lambda(r)\beta$ , so  $g = 0$ , an optimal ranking policy must always sort the pages by decreasing value of  $\tilde{R}_i$ .  $\square$*

The value of  $r$  obtained under this ordering, say  $r_0$ , is the maximal possible value, so we always have  $r \in [0, r_0]$ .

We conclude the section by offering an example that establishes that having a density for  $Y$  is not sufficient for the optimal policy to be deterministic and uniquely defined by  $\rho^*$ .

**Example 2.** Starting from Example 1, we add a third page with relevance  $R_3$  uniformly distributed over  $[0, \epsilon]$  for some small  $\epsilon > 0$ , and revenue  $G_3 = 0$ . We assume that  $\theta_3 = 1/4$ . Since  $R_3$  has a density,  $p(y) = 0$  for all  $y \in \Omega$ . For any  $\rho > 0$ , if  $\epsilon$  is small enough, this third page will always be ranked last, and its impact on  $h(r, g)$  is very small. Then the problem of ranking the first two pages becomes the same as Example 1, which means that the optimal policy must be randomized.

### 3.3 Some Illustrative Examples of our Methodology

This section provides two examples that, although simple and stylized, capture some features of the search market in the real world. The first example can model an SE that has content that competes with third-party CPs. Imagine that Google receives a video-search request and there are two pages that match the search; the first is from YouTube, owned by Google, while the second belongs to a competitor such as Dailymotion. Google's revenue generated by the YouTube page is positively correlated with the page relevance because of higher advertisement revenue and higher YouTube perception. Instead, Google's revenue generated by the Dailymotion page is negatively correlated with the page relevance because the more relevant, the more it diverts traffic from YouTube.

**Example 3.** Consider an instance with two pages where  $R_1$  and  $R_2$  are independent and uniformly distributed over  $[0, 1]$ ,  $G_1 = R_1$ ,  $G_2 = 1 - R_2$ , and  $\alpha_1 = \alpha_2 = \alpha$ . In addition, we let  $\lambda(r) = r$ ,  $\theta_1 = 1$ ,  $\theta_2 = 0$ , and  $\psi(R) = 1$ . We also define  $\tilde{\rho} := \alpha\rho$ .

At the optimal  $\rho$ , Page 1 will be ranked before Page 2 if and only if  $R_1 + \tilde{\rho}G_1 > R_2 + \tilde{\rho}G_2$ ; i.e., on the domain

$$D = \left\{ (R_1, R_2) : R_1 > \frac{\tilde{\rho}}{1 + \tilde{\rho}} + R_2 \frac{1 - \tilde{\rho}}{1 + \tilde{\rho}} \right\},$$

We define  $\bar{D} := [0, 1]^2 \setminus D$ . Then, since

$$0 < \frac{\tilde{\rho}}{1 + \tilde{\rho}} \leq \frac{\tilde{\rho}}{1 + \tilde{\rho}} + R_2 \frac{1 - \tilde{\rho}}{1 + \tilde{\rho}} \leq \frac{1}{1 + \tilde{\rho}} < 1,$$

we find that the LO- $\rho$  policy gives

$$\begin{aligned}
r &= \int_D r_1 dr_1 dr_2 + \int_{\bar{D}} r_2 dr_1 dr_2 \\
&= \int_0^1 \int_{\frac{\tilde{\rho}+r_2(1-\tilde{\rho})}{1+\tilde{\rho}}}^1 r_1 dr_1 dr_2 + \int_0^1 \int_0^{\frac{\tilde{\rho}+r_2(1-\tilde{\rho})}{1+\tilde{\rho}}} r_2 dr_1 dr_2 \\
&= \frac{1}{2} \int_0^1 \left( 1 - \left( \frac{\tilde{\rho} + r_2(1-\tilde{\rho})}{1+\tilde{\rho}} \right)^2 \right) dr_1 + \int_0^1 \frac{\tilde{\rho} + r_2(1-\tilde{\rho})}{1+\tilde{\rho}} r_2 dr_2 \\
&= \frac{1}{2} - \frac{1}{2(1+\tilde{\rho})^2} \frac{1-\tilde{\rho}^3}{3(1-\tilde{\rho})} + \frac{1}{1+\tilde{\rho}} \left( \frac{\tilde{\rho}}{2} + \frac{1-\tilde{\rho}}{3} \right) \\
&= \frac{1}{3} + \frac{\tilde{\rho}}{6(1+\tilde{\rho})^2} + \frac{1}{6} + \frac{1}{6(1+\tilde{\rho})} \\
&= \frac{2}{3} - \frac{\tilde{\rho}^2}{6(1+\tilde{\rho})^2}
\end{aligned}$$

and

$$\begin{aligned}
g/\alpha &= \int_D r_1 dr_1 dr_2 + \int_{\bar{D}} (1-r_2) dr_1 dr_2 \\
&= \frac{1}{3} + \frac{\tilde{\rho}}{6(1+\tilde{\rho})^2} + \int_{\bar{D}} dr_1 dr_2 - \frac{1}{6} - \frac{1}{6(1+\tilde{\rho})} \\
&= \frac{1}{3} + \frac{\tilde{\rho}}{6(1+\tilde{\rho})^2} + \int_0^1 \int_0^{\frac{\tilde{\rho}+r_2(1-\tilde{\rho})}{1+\tilde{\rho}}} dr_1 dr_2 - \frac{1}{6} - \frac{1}{6(1+\tilde{\rho})} \\
&= \frac{1}{3} + \frac{\tilde{\rho}}{6(1+\tilde{\rho})^2} + \int_0^1 \frac{\tilde{\rho} + r_2(1-\tilde{\rho})}{1+\tilde{\rho}} dr_2 - \frac{1}{6} - \frac{1}{6(1+\tilde{\rho})} \\
&= \frac{1}{3} + \frac{\tilde{\rho}}{6(1+\tilde{\rho})^2} + \frac{1}{2} - \frac{1}{6} - \frac{1}{6(1+\tilde{\rho})} \\
&= \frac{2}{3} - \frac{1}{6(1+\tilde{\rho})^2}.
\end{aligned}$$

The SE revenue is thus

$$\begin{aligned}
U_{SE} &= r(\beta + g) = \left( \frac{2}{3} - \frac{\tilde{\rho}^2}{6(1+\tilde{\rho})^2} \right) \left( \beta + \frac{2\alpha}{3} - \frac{\alpha}{6(1+\tilde{\rho})^2} \right) \\
&= \frac{\alpha\tilde{\rho}^2/36 - (3\beta + 2\alpha)\tilde{\rho}^2(1+\tilde{\rho})^2/18 - \alpha(1+\tilde{\rho})^2/9}{(1+\tilde{\rho})^4},
\end{aligned}$$

which we want to maximize over  $\tilde{\rho} \geq 0$ . Taking the derivative with respect to  $\tilde{\rho}$  and setting it to 0, we get the following equation, whose root divided by  $\alpha$  (since  $\tilde{\rho} = \alpha\rho$ ) provides us with  $\rho^*$ :

$$(3\beta + 2\alpha)\tilde{\rho}^3 + (6\beta + 5\alpha/2)\tilde{\rho}^2 + (3\beta - 5\alpha/2)\tilde{\rho} - 2\alpha = 0.$$

It follows from Proposition 3 (ii) that this  $\rho^*$  is unique. Moreover, Assumption A is satisfied, so  $\rho^*$  defines the order uniquely with probability 1.

To complete this example numerically, let us take  $\alpha = \beta = 1$ . Then,  $\rho^*$  is the unique positive root of  $5\rho^3 + 17\rho^2/2 + \rho/2 - 2$ , which is  $\rho^* \approx 0.412149553$ . Figure 1 shows the expected SE revenue




Fig. 1: Expected SE Revenue in terms of  $\rho$  when  $\alpha = \beta = 1$

for this example, as a function of  $\rho$ . One can also compute that for  $\rho = \rho^*$ , we have  $r = 0.6524696521$  and  $g = 0.583089554$ , and then  $h(r, g) = r/(1 + g) = 0.412149553 = \rho^*$ , as expected.  $\square$

In our next example, we consider payments that are all-or-nothing, where the SE gets a revenue when it is also the CP serving the corresponding page. Otherwise, the SE does not have any financial gain when showing the link to the page. An alternative interpretation of this example is that content, instead of being served by the SE, is served by CPs some of whom agree to pay the SE a fixed price, normalized to 1, for each click to their pages served from the SE's output. That price is not tied to a fixed position for the link; it just provides an incentive so the SE favors links with  $G_i = 1$  in its ranking.

**Example 4.** Consider an instance with two pages ( $M = 2$ ) where revenues  $G_i$  can only take values 0 and 1. Indeed, for  $i = 1, 2$ ,  $G_i$  is a Bernoulli random variable with parameter  $p$ ,  $R_i$  has a uniform distribution over  $[0, 1]$  independent of  $G_i$ , and  $\alpha_i = 1$ . We let  $(\theta_1, \theta_2) = (1, 0)$ ,  $\lambda(r) = r$ , and  $\psi(R) = 1$ . Note that this is equivalent to saying that only one page is displayed in the search output.

The density of  $Y$  is a mixture of two uniforms, which verifies Assumption A. Focusing on LO- $\rho$  policies, we derive explicit formulas for  $r = r(\rho)$ ,  $g = g(\rho)$ , and  $\varphi(r(\rho), g(\rho))$ . The fixed point  $\rho^*$  can be computed from them. To start, we compute the average relevance  $r = r(\rho)$ . We distinguish two cases for the vector  $((R_1, G_1), (R_2, G_2))$ :

1. If  $G_1 = G_2$ , only the most relevant link is displayed, resulting in conditional expected relevance  $\mathbb{E}[\max(R_1, R_2) \mid G_1 = G_2] = 2/3$ .
2. If  $G_1 \neq G_2$ , we can assume (possibly by swapping the roles of pages 1 and 2) that  $G_1 = 1$  and  $G_2 = 0$ . If  $R_1 + \rho \geq R_2$ , link 1 is displayed and the observed relevance is  $R_1$ ; otherwise, the observed relevance is  $R_2$ . Note that if  $\rho > 1$ , link 1 is always shown, leading to an expected observed relevance of  $1/2$ . If  $\rho \leq 1$ , the expected relevance conditional on  $(G_1, G_2)$  is

$$\mathbb{E}[R_1 \mathbb{1}_{\{R_1 + \rho > R_2\}} + R_2 \mathbb{1}_{\{R_1 + \rho \leq R_2\}} \mid G_1 = 1, G_2 = 0]$$

$$\begin{aligned}
&= \int_{r_1=0}^1 \int_{r_2=0}^1 r_1 \mathbb{1}_{\{r_1+\rho>r_2\}} + r_2 \mathbb{1}_{\{r_1+\rho\leq r_2\}} dr_2 dr_1 \\
&= \int_{r_1=0}^1 \left( \int_{r_2=0}^{\min(1,r_1+\rho)} r_1 dr_2 + \int_{r_2=\min(1,r_1+\rho)}^1 r_2 dr_2 \right) dr_1 \\
&= \int_{r_1=0}^1 \left( r_1 \min(1, r_1 + \rho) + \frac{1 - \min(1, (r_1 + \rho)^2)}{2} \right) dr_1 \\
&= \int_{r_1=0}^{1-\rho} \left( r_1(r_1 + \rho) + \frac{1 - (r_1 + \rho)^2}{2} \right) dr_1 + \int_{r_1=1-\rho}^1 r_1 dr_1 \\
&= \int_{x_1=\rho}^1 \left( (x_1 - \rho)x_1 + \frac{1 - x_1^2}{2} \right) dx_1 + \int_{r_1=1-\rho}^1 r_1 dr_1 \\
&= \int_{x_1=\rho}^1 \left( \frac{1}{2} - \rho x_1 + \frac{x_1^2}{2} \right) dx_1 + \frac{1 - (1 - \rho)^2}{2} \\
&= \frac{2}{3} - \frac{\rho^2}{2} + \frac{\rho^3}{3}.
\end{aligned}$$

Combining the four possibilities for  $(G_1, G_2)$ , the overall expected relevance for the LO- $\rho$  policy is

$$r = r(\rho) = \frac{2}{3} + p(1-p)\bar{\rho}^2 \left( \frac{2\bar{\rho}}{3} - 1 \right), \quad (16)$$

where  $\bar{\rho} := \min(1, \rho)$ .

Similarly, to compute the expected revenue  $g = g(\rho)$  per request, we consider two cases:

1. If  $G_1 = G_2$ , the expected revenue is 0 if  $G_1 = 0$ , and 1 otherwise.
2. If  $G_1 \neq G_2$ , we can assume again that  $G_1 = 1$  and  $G_2 = 0$ . Again, if  $\rho > 1$ , link 1 is always shown and the revenue is 1. If  $\rho \leq 1$ , the expected revenue conditional on  $(G_1, G_2)$  is

$$\begin{aligned}
\mathbb{E}[\mathbb{1}_{\{R_1+\rho>R_2\}} \mid G_1 = 1, G_2 = 0] &= \int_{r_1=0}^1 \int_{r_2=0}^1 \mathbb{1}_{\{r_1+\rho>r_2\}} dr_2 dr_1 \\
&= \int_{r_1=0}^1 \min(1, r_1 + \rho) dr_1 \\
&= \int_{r_1=0}^{1-\rho} (r_1 + \rho) dr_1 + \int_{r_1=1-\rho}^1 1 dr_1 \\
&= \int_{x_1=\rho}^1 x_1 dx_1 + \rho \\
&= 1 - \frac{(1 - \rho)^2}{2}.
\end{aligned}$$

Regrouping all cases, we obtain

$$g = g(\rho) = p^2 + 2p(1-p) \left( 1 - \frac{(1 - \bar{\rho})^2}{2} \right). \quad (17)$$



Note that both  $r(\rho)$  and  $g(\rho)$  are constant for  $\rho \geq 1$ , so we can reduce the search for an optimal  $\rho$  to the interval  $[0, 1]$ , and in that interval  $\bar{\rho} = \rho$ . With  $\lambda(r) = r$ , the expected revenue per unit of time is  $U_{SE}(\rho) = r(\rho) \cdot (\beta + g(\rho))$ , which equals

$$\left( \frac{2}{3} + p(1-p)\rho^2 \left( \frac{2\rho}{3} - 1 \right) \right) \cdot \left( \beta + \left( p^2 + 2p(1-p) \left( 1 - \frac{(1-\rho)^2}{2} \right) \right) \right).$$

Figure 2 depicts the expected revenue as a function of  $\rho$ , along with  $r(\rho)$  and  $g(\rho)$ , for  $\beta = 1$  and  $p = 1/2$ . While  $g(\rho)$  increases and  $r(\rho)$  decreases with  $\rho$ , the maximal revenue is obtained by taking  $\rho$  around 0.4. This optimal  $\rho$  uniquely determines the optimal policy (with probability 1).  $\square$




Fig. 2: Expected SE revenue per unit time for  $\beta = 1$  and  $p = 1/2$

The previous example illustrates that to appropriately solve the tradeoff between short-term revenue coming from payments and long-term revenue coming from more exposure due to higher relevance, one must place an appropriate weight on short-term revenues and on relevance (the former being around 40% of the latter in our example). This will provide short-term benefits to the SE without impairing its possibility to attract future users.

#### 4 Finding Optimal Rankings by Computing $\rho^*$

In this section, we discuss how to find the optimal  $\rho^*$  that allows the SE to determine the revenue-maximizing ranking easily. Proposition 1 shows that to achieve an optimal revenue, the SE should rank the items in the request  $Y$  by decreasing order of  $\tilde{R}_i + \rho\tilde{G}_i$ , for a properly chosen  $\rho = h(r, g)$ . But  $\rho$  depends on  $r$  and  $g$ , which in turn depend on the selected policy  $\mu$  and are unknown a-priori. Moreover, typically, this dependence is not expressed in a closed-form formula. In the examples of Section 3.3, we were able to derive explicit analytical expressions for  $r(\rho)$  and  $g(\rho)$ , and use them to find the optimal  $\rho$ . Unfortunately, instances of real size do not admit such closed-form derivations

and they would usually have to be estimated through simulation. This motivates the following *stochastic root-finding* problem: estimate a root of  $\tilde{h}(\rho) - \rho = 0$  when only noisy estimates of  $\tilde{h}$  can be obtained, via simulation. Several algorithms have been designed and studied for this type of problem; see, e.g., [23] and the references therein. We assume that a root exists and is unique.

An estimator  $\hat{h}_n(\rho)$  of  $\tilde{h}(\rho)$  at any given value of  $\rho$  can be defined and computed as follows. We generate  $n$  independent realizations  $Y_1, \dots, Y_n$  of  $Y$ , with  $Y_i = (M_i, R_{i,1}, G_{i,1}, \alpha_{i,1}, \dots, R_{i,M_i}, G_{i,M_i}, \alpha_{i,M_i})$ . For each  $i$ , we order the triples  $(R_{i,k}, G_{i,k}, \alpha_{i,k})$  by decreasing order of  $\tilde{R}_{i,k} + \rho \tilde{G}_{i,k}$ , and we compute  $r_i(\rho) = \sum_{k=1}^{M_i} \theta_k \tilde{R}_{i,\delta(k)}$  and  $g_i(\rho) = \sum_{k=1}^{M_i} \theta_k \tilde{G}_{i,\delta(k)}$ . Unbiased estimators of  $r(\rho)$  and  $g(\rho)$  are then  $\hat{r}_n(\rho) = (1/n) \sum_{i=1}^n r_i(\rho)$  and  $\hat{g}_n(\rho) = (1/n) \sum_{i=1}^n g_i(\rho)$ , respectively. They lead to the estimator

$$\hat{h}_n(\rho) = \varphi(\hat{r}_n(\rho), \hat{g}_n(\rho)), \quad (18)$$

which is generally biased for finite  $n$  when  $\varphi$  is nonlinear, but is consistent, and the bias typically decreases as  $\mathcal{O}(1/n)$  [2]. For the special case where  $\varphi(r, g) = \lambda(r)(\beta + g)$ , this gives  $\hat{h}_n(\rho) = \lambda(\hat{r}_n(\rho))(\beta + \hat{g}_n(\rho))$ . A confidence interval for  $\tilde{h}(\rho)$  can be computed using the Delta method [2], under the assumption that  $\hat{r}_n(\rho)$  and  $\hat{g}_n(\rho)$  have (approximately) a normal distribution.

When searching for a root of  $\tilde{h}(\rho) - \rho$ , or if we want to estimate the function  $\tilde{h}$  over some interval, we need to compute  $\hat{h}_n(\rho)$  at many values of  $\rho$ . This can be done using *common random numbers* (CRN), which means that we use exactly the same  $n$  realizations  $Y_1, \dots, Y_n$  at all values of  $\rho$  at which we perform a function evaluation, or using *independent random numbers* (IRN), in which case we draw a fresh independent sample  $Y_1, \dots, Y_n$  at each  $\rho$  where we estimate  $\tilde{h}(\rho)$ . In the CRN case,  $\hat{h}_n(\rho)$  becomes a deterministic function of  $\rho$  and this function typically varies much less than in the IRN case. The *sample average optimization* method consists in optimizing this sample function  $\hat{h}_n(\rho)$  defined with CRNs. However, for any fixed  $n$ , this sample function is piecewise-constant in  $\rho$ , because it depends on  $\rho$  only via the selected permutation for each  $i$ , and therefore only takes a finite number of values as a function of  $\rho$ . As a result, its derivative is zero almost everywhere and (in general)  $\hat{h}_n(\rho) - \rho$  has no exact root. Therefore, the best we can do for fixed  $n$  is to compute an *approximate* root  $\hat{\rho}_n^*$  of  $\hat{h}_n(\rho) - \rho$ , and for this, any method that relies on the derivative of  $\hat{h}_n(\rho)$  must be ruled out. We can compute the approximate root either by a method that does not rely on derivatives (such as binary search), or by a derivative-based method (e.g., a Newton-type method) by approximating the derivative with finite differences. Thus, we can compute  $\hat{\rho}_n^*$  such that  $\epsilon_n = |\hat{h}_n(\hat{\rho}_n^*) - \hat{\rho}_n^*|$  is small, and do this for an increasing sequence of values of  $n$ , in a way that  $\epsilon_n \rightarrow 0$  when  $n \rightarrow \infty$ . This is possible under the assumption that  $\hat{h}_n \rightarrow \tilde{h}$  uniformly when  $n \rightarrow \infty$ , which usually occurs with CRNs (under mild conditions). For each considered sample size  $n$ , we would use the approximate root  $\hat{\rho}_n^*$  as a starting point when finding the approximate root for the next (larger) value of  $n$ .

Another approach is to use a Robbins-Monro-type *stochastic approximation* (SA) iterative method; see [23] for an overview and convergence results. For the situation where  $\tilde{h}(\rho) - \rho$  is decreasing in  $\rho$ , SA starts from some  $\rho_0$  and generates iterates of the form

$$\rho_{j+1} = \rho_j + a_j(\hat{h}_{n_j}(\rho_j) - \rho_j), \quad (19)$$

where  $\hat{h}_{n_j}(\rho_j)$  is an estimate of  $\tilde{h}(\rho_j)$  based on sample size  $n_j$ . These estimates are independent across values of  $j$ , and  $\{a_j, j \geq 0\}$  is a slowly-decreasing sequence such that  $\sum_{j=0}^{\infty} a_j = \infty$  and  $\sum_{j=0}^{\infty} a_j^2 < \infty$ . Note that there is no need to have  $n_j \rightarrow \infty$ ; one can take  $n_j$  as a small constant independent of  $j$ . If we replace  $a_j$  by the inverse derivative  $1/(\tilde{h}'(\rho_j) - 1)$  and the estimate of  $\tilde{h}(\rho_j)$

by its exact value, we obtain the Newton method, which usually converges much faster, but requires knowledge of the function and of its derivative (or accurate estimators and  $n_j \rightarrow \infty$ ), in contrast to SA. On the other hand, without a good choice of the  $a_j$ 's, SA might converge extremely slowly.

If we replace  $a_j$  by 1 in (19), we obtain

$$\rho_{j+1} = \rho_j + (\hat{h}_{n_j}(\rho_j) - \rho_j) = \hat{h}_{n_j}(\rho_j). \quad (20)$$

If  $n_j \rightarrow \infty$ , this iteration becomes equivalent in the limit to the mapping  $\rho \rightarrow \tilde{h}(\rho)$ . Recall that  $\rho \rightarrow \tilde{h}(\rho)$  is a *contraction mapping* if there is a constant  $\gamma \in [0, 1)$  such that

$$|\tilde{h}(\rho) - \tilde{h}(\rho')| \leq \gamma |\rho - \rho'|$$

for all  $\rho, \rho' \geq 0$ . A sufficient condition for this to hold is that  $|\tilde{h}'(\rho)| \leq \gamma$  for all  $\rho$  (in the region of interest). When this holds, we can start from some  $\rho_0 > 0$  and iterate:  $\rho_{j+1} = \tilde{h}(\rho_j)$ , for  $j = 1, 2, \dots$ . Then, the fixed-point theorem for contraction mappings [6] guarantees that  $\rho_j \rightarrow \rho^*$  at a geometric rate:  $|\rho_j - \rho^*| \leq \gamma^j |\rho_0 - \rho^*|$ , which provides very fast convergence when  $\gamma \ll 1$ . In practice, we can replace  $\tilde{h}(\rho_j)$  by  $\hat{h}_{n_j}(\rho_j)$ , and convergence to  $\rho^*$  will occur if  $n_j \rightarrow \infty$  when  $j \rightarrow \infty$ . On the other hand, if  $n_j$  does not increase with  $j$ ,  $\rho_j$  will generally not converge to  $\theta^*$ . If  $n_j$  is fixed to some large constant  $n$  and we use IRN,  $\rho_j$  will never converge but wander around in a small neighborhood of  $\theta^*$ . If we use CRNs, it will converge to a value close to  $\theta^*$ , but generally different.

It is very common in our model that  $\rho \rightarrow \tilde{h}(\rho)$  is a contraction mapping. In particular, this holds in all the examples considered in this paper. Generating iterates of (20), we verified that in all cases it converged very quickly to a very good approximation of  $\rho^*$ .

To illustrate the previous discussion about computation we solve the examples provided in Section 3.3 numerically.

**Example 5.** We revisit Example 3 and estimate  $r(\rho)$ ,  $g(\rho)$ , and  $U_{SE}(\rho) = \varphi(r(\rho), g(\rho))$  with IRN. In particular, we take a sample of size  $n = 10^7$ , and  $\rho$  from 0 to 1 with a step size of 0.001. The plot on the left of Figure 3 displays the estimates of  $U_{SE}(\rho)$ . This gives an idea of the high-frequency estimation noise achieved with IRN. The true maximum of  $U_{SE}(\rho)$  is found at  $\rho^* = 0.41214955$  whereas the numerical estimate is  $\rho = 0.437$ . We see that, even with this large sample size, the noise is significant compared with the variation of  $U_{SE}(\rho)$  around  $\rho^*$ . This illustrates the fact that sample-average optimization with IRN is not a good method to approximate the optimal ranking policy, because of the large high-frequency noise in the sample function.

We also applied the mapping (20) for several iterations, starting at  $\rho_0 = 0$ , with a fixed sample size of  $n_j = 10^7$  for all  $j$ , and IRN across iterations. This gave already  $\rho_1 = 0.44446$  at the first iteration and  $\rho_4 = 0.4121$  (accurate up to four digits) after four iterations. Thus, this method gets close to the optimum very quickly. For comparison, we ran the same method with CRN. The values of  $\rho_j$  with both methods are shown in Table 1, for  $j = 1, \dots, 6$ . With both methods,  $\rho_j$  provides a good approximation of  $\rho^*$  very quickly. With CRN, it converges to 0.4121425 for  $j \geq 7$ , which is not the correct value but is accurate to five digits.

To show that we indeed have a contraction mapping, recall that for this example,  $\lambda(r) = r$ ,  $r = 2/3 - \rho^2/(6(1 + \rho)^2)$ , and  $g = 2/3 - 1/(6(1 + \rho)^2)$ . This gives

$$\tilde{h}(\rho) = \frac{2/3 - \frac{\rho^2}{6(1+\rho)^2}}{\beta + 2/3 - \frac{1}{6(1+\rho)^2}} \quad \text{and} \quad \tilde{h}'(\rho) = -\frac{2(4\rho^2 + 6\beta\rho^2 + 6\beta\rho + 7\rho + 4)}{(6\beta + 12\beta\rho + 6\beta\rho^2 + 3 + 8\rho + 4\rho^2)^2}.$$

Method	$\rho_1$	$\rho_2$	$\rho_3$	$\rho_4$	$\rho_5$	$\rho_6$
IRN	0.4444478	0.4102066	0.4122955	0.4121269	0.4121527	0.4121252
CRN	0.4444478	0.4101850	0.4122638	0.4121351	0.4121428	0.4121424

Tab. 1: Values of  $\rho_j$  at the first six iterations of (20) for Example 5, with IRN and CRN

Method	$\rho_1$	$\rho_2$	$\rho_3$	$\rho_4$	$\rho_5$	$\rho_6$
IRN	0.4444471	0.377720	0.387115	0.3857771	0.3860725	0.3859246
CRN	0.4444471	0.377670	0.387079	0.3857318	0.3859223	0.3858940

Tab. 2: Values of  $\rho_j$  at the first six iterations of (20) for Example 6, with IRN and CRN

For  $\beta = 1$ , one can verify that  $\tilde{h}'(\rho)$  is negative and increasing, with  $|\tilde{h}'(\rho)| \leq |\tilde{h}'(0)| = 8/81 < 1$ . Therefore the mapping  $\rho \rightarrow \tilde{h}(\rho)$  is contracting with  $\gamma = 8/81$ .  $\square$

**Example 6.** We also revisit Example 4 and solve the problem numerically. The plot on the right of Figure 3 shows the estimates of  $r(\rho)$ ,  $g(\rho)$ , and  $U_{SE}(\rho)$ , computed with IRN with  $n = 10^5$ . We also superimpose the corresponding exact curves. Again, we applied (20) for six iterations, starting with  $\rho_0 = 0$ , and a fixed sample size of  $n_j = 10^7$  for all  $j$ . The results are in Table 2. We find that  $\rho^* \approx 0.3859$ .

For this example, with the expressions previously derived for  $r$  and  $g$ , we get

$$\tilde{h}(\rho) = \frac{2/3 + p(1-p)\bar{\rho}^2(2\bar{\rho}/3 - 1)}{\beta + p^2 + 2p(1-p)(1 - (1-\bar{\rho})^2/2)}$$

and

$$\tilde{h}'(\rho) = -\frac{2}{3}p(1-p)(1-\bar{\rho}) \frac{(3\bar{\rho}\beta + 3p\bar{\rho} + 3p\bar{\rho}^2 - p\bar{\rho}^3 - 3p^2\bar{\rho}^2 + p^2\bar{\rho}^3 + 2)}{(\beta + p + 2p\bar{\rho} - p\bar{\rho}^2 - 2p^2\bar{\rho} + p^2\bar{\rho}^2)^2}.$$

For  $\beta = 1$  and  $p = 1/2$ , one can verify numerically that for  $0 \leq \rho \leq 1$ ,  $\tilde{h}'(\rho)$  is negative and achieves a maximum absolute value of approximately  $0.15 < 1$  (although the derivative is not monotone). Hence, we have a contraction mapping with  $\gamma \approx 0.15$  in that area.  $\square$

## 5 Comparison of the Neutral and Non-Neutral Ranking Policies

In this section, we show via numerical examples how the theory developed earlier can be used to study the impact of different ranking policies on various performance indicators such as consumer welfare (captured by expected relevance), SE and CP revenue. In particular, we compare neutral ranking policies, where  $\rho = 0$ , with non-neutral ones, where the SE chooses the optimal  $\rho^*$ .

### 5.1 A Vertically Integrated SE with a CP

**Example 7.** We first focus on a specific type of request which can be served by either third-party CPs or by the SE itself. This is typical for many search categories where the SE also provides content (e.g., video, weather, finance, news, maps, flight information, and so on). In this case, a limited number of CPs compete with the SE, and the parameters  $r$ ,  $g$ , and  $\lambda(r)$  for the instance correspond to just this type of request. Let us assume that always ten pages match a request ( $M = 10$ ). Nine




Fig. 3: (Left) Estimate of expected SE revenue per unit time in terms of  $\rho$  for  $\alpha = \beta = 1$  in Example 3, and (Right) Estimate of expected SE revenue per unit time for  $\alpha = 1$ ,  $\beta = 1$ ,  $p = 1/2$  in Example 4

$\theta_1$	$\theta_2$	$\theta_3$	$\theta_4$	$\theta_5$	$\theta_6$	$\theta_7$	$\theta_8$	$\theta_9$	$\theta_{10}$
0.364	0.125	0.095	0.079	0.061	0.041	0.038	0.035	0.03	0.022

Tab. 3: CTR values used in the simulations of Section 5

of those pages are served by third-party CPs but one of them is served by the SE directly. Perhaps renumbering CPs, we have that  $\alpha_1 = 1$ , and  $\alpha_2 = \dots = \alpha_{10} = 0$ . In addition to the revenue coming from Page 1, the SE also receives an expected revenue of  $\beta = 1$  per request from sponsored links. For  $i = 1, \dots, 10$ ,  $R_i$  and  $G_i$  are all independent random variables uniformly distributed over  $[0, 1]$ , and  $\text{CTR}(i) = \theta_i$  as specified in Table 3. Those numbers were taken from the first table in [11], which contains the observed relative numbers of clicks according to the position: the actual CTRs should therefore be proportional to those numbers, and the value of the multiplicative constant has no impact on our derivations (hence we take it equal to 1). Finally, we set  $\lambda(r) = r$ , and  $\psi$  to be the unit function.

The  $M$  pages are ranked by the SE by decreasing value of  $\tilde{R}_i + \rho \tilde{G}_i$ , for the correct constant  $\rho \geq 0$ . Note that for  $i > 1$ ,  $\tilde{G}_i = 0$  because  $\alpha_i = 0$ . To illustrate the dependence on  $\rho$ , Figure 4 shows the SE revenue  $U_{\text{SE}}(\rho)$ , as  $\rho$  varies, as well as the relevance  $r(\rho)$ , the revenue and the visit rate for CP 1 and for third-party CPs. All revenues are expressed as values per time unit. As discussed earlier, the more  $\rho$  increases, the more the SE favors CP 1, decreasing the overall relevance and increasing the visit rate to CP 1. The trade-off between short-term revenue and number of visits tells the SE to choose  $\rho^* \approx 0.55$ . Note that the bias affects only CP 1 and that the relative positions of all other CPs remain the same as in the neutral ranking. Consequently, the relevance  $r(\rho)$  is only marginally affected by  $\rho$  in this case. If  $R_1$  was stochastically much smaller than the other  $R_i$ 's (e.g., uniform over  $[0, \epsilon]$  for a small  $\epsilon$ ), then the impact of  $\rho$  would be larger. When  $\rho \rightarrow \infty$ , CP 1 is always ranked first, so the relevance  $r(\rho)$  becomes

$$r(\infty) = \left( \frac{\theta_1}{2} + \sum_{i=1}^9 \theta_{i+1} \mathbb{E}[U_{(10-i)}] \right) = \frac{\theta_1}{2} + \sum_{i=1}^9 \theta_{i+1} \frac{(10-i)}{10} \approx 0.517,$$

where  $U_{(1)}, \dots, U_{(9)}$  are independent random variables uniformly distributed over  $[0, 1]$  sorted by increasing order (the order statistics), and the CP 1 visit rate is  $\theta_1 r(\infty) \approx 0.188$ .

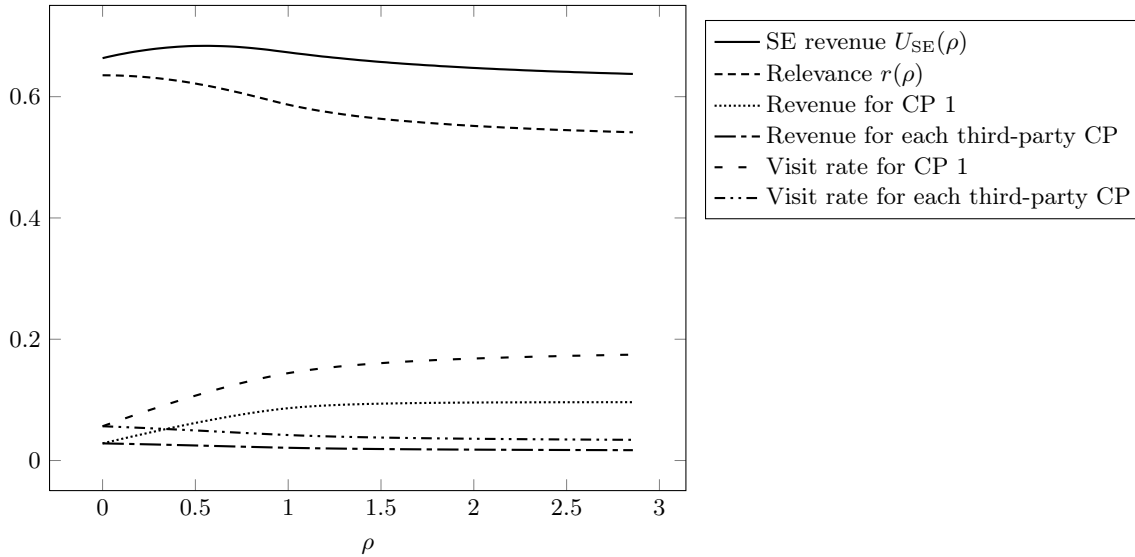


Fig. 4: Performance measures as a function of  $\rho$  (simulation results)

To assess the sensitivity of the SE strategy to advertising, we now examine how results change for different values of  $\beta$ , i.e., depending on the level of advertisement revenues. This shows the tradeoff that the SE faces for different types of requests. For search keywords related to, e.g., airline tickets, hotel reservations, or retailer products, the SE may expect to make more profit by showing its own content among organic links than through sponsored search because requests of this kind may produce conversions, whereas for keywords that are appealing in the sponsored search market the SE may try to make the search as relevant as possible to boost that revenue stream. Figure 5 plots  $\rho^*$  as  $\beta$  varies while Figure 6 plots the ensuing revenue for CP 1 and for each third-party CP. The curves shown in the figures were estimated by simulation, using the iterative fixed-point method for  $\rho^*$ , with a fixed sample size of  $n = 10^7$  at each step. When  $\beta$  grows,  $\rho^*$  tends to zero, because the revenue from sponsored links dominates, making it rewarding for the SE to improve quality to attract more users. In conclusion, the impact of non-neutrality is small because biasing the ranking only attracts limited additional revenue. Instead, when  $\beta$  is small, sponsored links do not pay off and it becomes worthwhile for the SE to sacrifice relevance to some extent to boost revenue from gains of CP 1. In the extreme case when  $\beta = 0$ , we have  $\rho^* = \infty$ , so CP 1 is always placed at the top regardless and the other CPs are sorted by decreasing order of relevance. This gives an average revenue of 0.09619 for CP 1 and 0.01695 for any other CP (even though all CPs have the same relevance and gain distributions). Although not shown in the figure, we remark that  $U_{SE}$  tends to grow linearly with  $\beta$ , which means that the increasing revenues of sponsored search dominate the possible revenue coming from CP 1. To illustrate the impact of non-neutrality, Table 4 reports the variations of the most relevant performance metrics when  $\rho = \rho^*$  is used instead of  $\rho = 0$  (neutral ranking), for different values of  $\beta$ . The table illustrates that while the impact on the perceived quality (relevance) remains small (around 10%), the impact on the visibility and the revenues of the SE-owned CP is substantial: by being non-neutral, the SE can multiply the revenues of its CP




Fig. 5: Value of  $\rho^*$  as a function of  $\beta$  (simulation results)




Fig. 6: CP revenues per time unit as a function of  $\beta$  (simulation results)

	Relevance	CP 1 revenue	other CP revenue	CP 1 visit rate	other CP visit rate
Neutral, $\rho = 0$ (reference case optimal for $\beta = \infty$ )	0.635	0.028	0.0283	0.057	0.057
Non-neutral, $\rho = 0.559$ (optimal for $\beta = 1$ )	0.618 (-3%)	0.066 (+136%)	0.0243 (-14%)	0.112 (+96%)	0.049 (-14%)
Non-neutral, $\rho = 0.924$ (optimal for $\beta = .5$ )	0.592 (-7%)	0.084 (+200%)	0.0215 (-24%)	0.140 (+146%)	0.043 (-25%)
Non-neutral, $\rho = 1.374$ (optimal for $\beta = .25$ )	0.568 (-11%)	0.093 (+232%)	0.0193 (-32%)	0.158 (+177%)	0.039 (-32%)

Tab. 4: Impacts of a non-neutral ranking for the scenario of Section 5.1

by a factor 2.8 and its visit rate by more than a factor of 3. On the other hand, the other CPs see their revenues and visit rates reduced by 14% to 32%, a significant loss that is likely to affect their possibilities of being profitable in the long term.

Finally, we explore the sensitivity of outcomes to the number of available results. Figure 7 plots  $\rho^*$  as a function of  $M$  while Figure 8 plots revenues as a function of the number of matching pages  $M$ . We include curves for both the neutral ( $\rho = 0$ ) and non-neutral ( $\rho = \rho^*$ ) regimes to compare both situations. As before, we estimate these values using the fixed-point algorithm with  $n = 10^7$  at each step. As  $M$  increases,  $\rho^*$  increases too: The SE can give more weight to CP 1 and increase its revenue while making less damage to the relevance, because placing CP 1 higher has less impact on the overall relevance when  $M$  is larger. As a result, the revenue of CP 1 when  $\rho = \rho^*$  increases with  $M$ , and so does the advantage of CP 1 over the other CPs. The loss of revenue of the other CPs seems close to constant as a function of  $M$ .

## 5.2 Vertical Integration and Investment

**Example 8.** Continuing with the example of vertical integration, we now assume that one of the nine third-party CPs, say CP 2, invests in quality and manages to improve the relevance distribution. More specifically, we assume that when it invests  $z > 0$ , the relevance of CP becomes uniformly distributed over  $[0, 1+20z]$  (instead of over  $[0, 1]$ ). The other parameters and distributions, including the distribution of its gain  $G_2$ , are unchanged. Figures 9 and 10 show simulation results when the SE ranks CPs according to  $\tilde{R}_i + \rho\tilde{G}_i$ , for varying values of  $\rho$ , and when  $z = 2$ . For a neutral ranking ( $\rho = 0$ ), CP 2 logically makes more revenue than the other CPs, since it regularly gets higher ranking. However, when  $\rho$  increases and exceeds about 0.8, CP 1 becomes the one with highest revenue, despite its (stochastically) lower relevance.

We now take the perspective of CP 2, and compute its optimal decision. CP 2 invests  $z$  in quality to modify its relevance distribution to  $[0, 1+20z]$ , anticipating that the SE is going to rank requests according to  $\rho^*$ . (We assume that the SE can learn the distribution of relevance of all CPs quickly.) Therefore, CP 2's profit equals the revenue from the search market minus  $z$ . To find the optimal value of  $z$  we simulated the outcomes for  $z \in [0, 0.45]$ . Figures 11 and 12 plot the resulting curves. In both figures, we see that differences between neutral and non-neutral revenues






Fig. 7: Value of  $\rho^*$  as a function of  $M$  (simulation results)

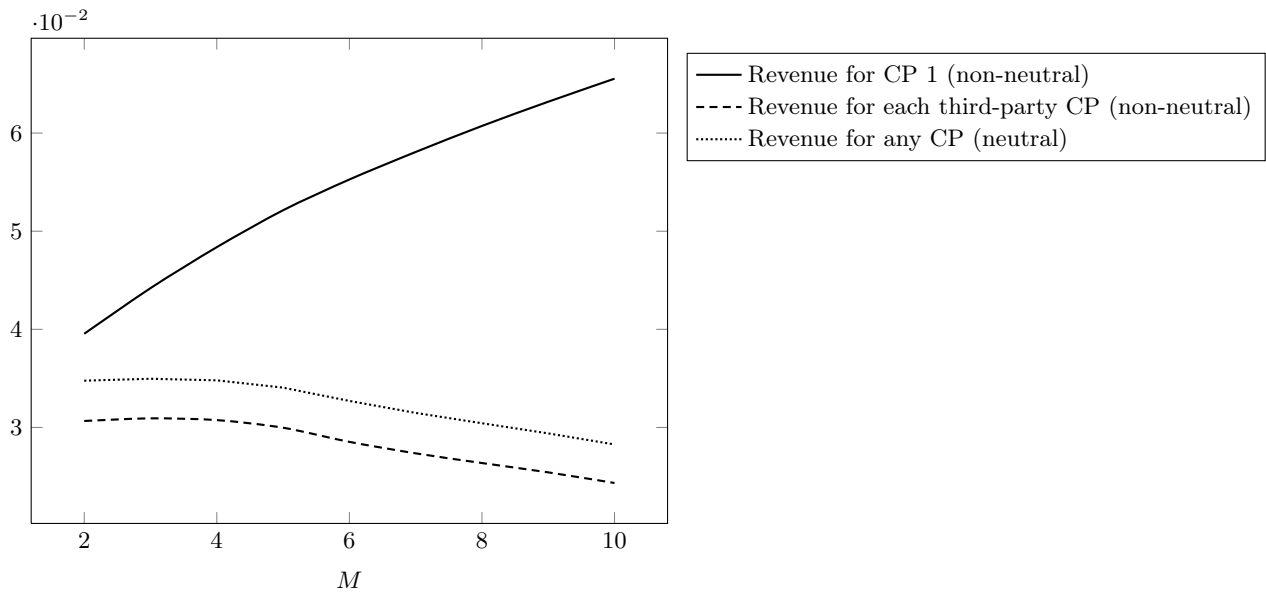


Fig. 8: CP revenues per time unit as a function of  $M$  (simulation results)




Fig. 9: Relevance and revenues per time unit for the case of vertical integration with investment




Fig. 10: Number of visits per time unit for the case of vertical integration with investment

are small, except for CP 1. This is particularly true for CP 2. This means that, at least in this case, non-neutrality does not deter innovation. Actually, the optimal investment level under both regimes coincide and is equal to  $z^* = 0.025$ . Optimal profits, though, vary. They are 0.037 for the neutral case and 0.0296 for the non-neutral one; see Figure 13 where we show CP 2 profits as a function of the investment  $z$ .

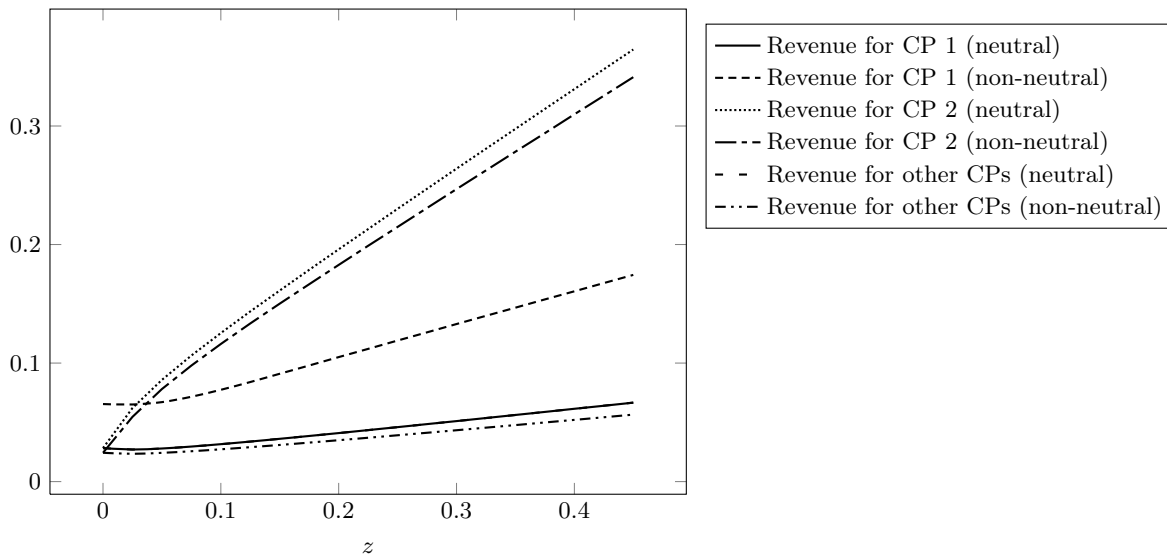


Fig. 11: Revenues per time unit as a function of CP 2 investment

## 6 Conclusion

We have introduced a new modeling framework that allows online platforms to rank items accounting for both short-term and long-term revenues. The long-term impact is captured by the arrival rate of requests, which is an increasing function of the average relevance of displayed results. Under appropriate regularity conditions, we proved that although we have to choose an ordering among an exponential number of possibilities and the objective function is nonlinear, the task reduces to computing a linear combination between relevance and short-term profits for each item and then sorting items with respect to those numbers. Henceforth, the whole problem reduces to finding the appropriate constant used in the linear combination. We have also provided algorithms to find such constant.

Our results might prove useful to platform owners (search engines, classified ads websites, online retailers) to navigate the tradeoff between short-term and long-term effects when defining their ranking strategies. They can also be of interest to regulators, seeking to understand the behavior of revenue-oriented platforms and to anticipate the impact of regulatory interventions, which is of particular importance with regard to the current *search neutrality* debate.

Future work will take several directions. In particular, we plan to (i) work on the design of optimal randomized policies as highlighted in the case of discrete distributions of requests; this will also allow us to relax Assumption A in the case of a continuous distribution for requests; (ii) gather

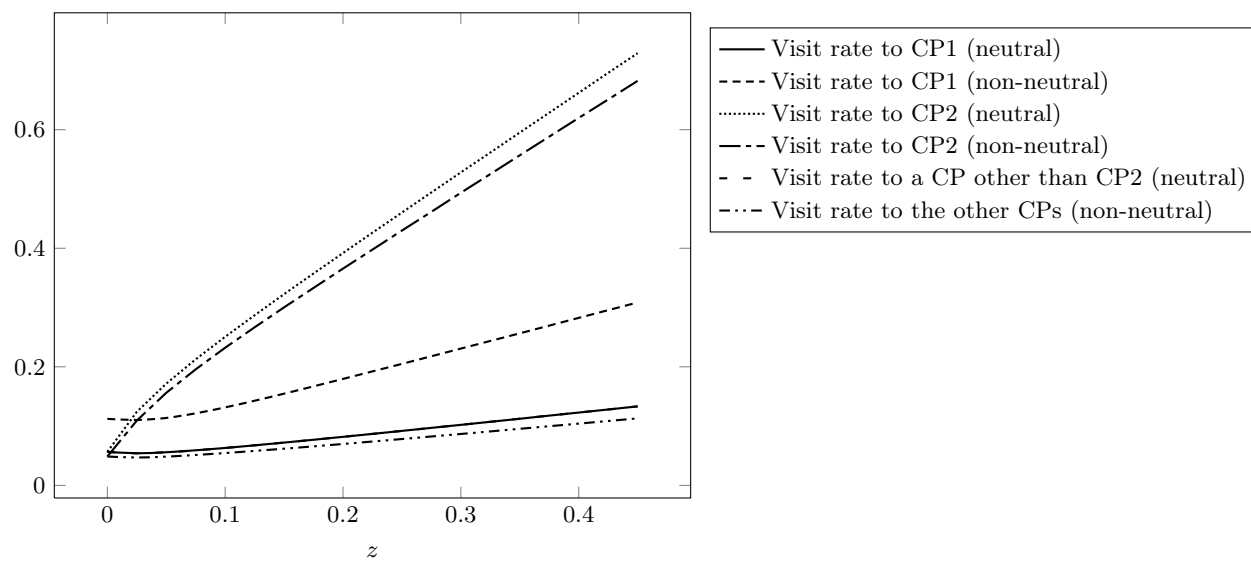


Fig. 12: Visit rates to various CPs as a function of CP 2 investment




Fig. 13: Profit per time unit as a function of CP 2 investment

real data and providing a practical case study; (iii) perform a profound study of the implications of profit-maximizing platforms on the online economy to shed light on the search-neutrality debate.

## Acknowledgement

This work has been supported by Inria's associated team MOCQUASIN, as well as an Inria International Chair, a Canada Research Chair, an NSERC Discovery Grant to P. L'Ecuyer and by CONICET Argentina Grant PIP 112-201201-00450CO and FONCYT Argentina Grant PICT 2012-1324.

## References

- [1] E. Altman, A. Legout, and Y. Xu. Network non-neutrality debate: An economic analysis. In *Proc. of Networking*, pages 68–81, Barcelona, Spain, May 2011. LNCS 6641, Springer Verlag.
- [2] S. Asmussen and P. W. Glynn. *Stochastic Simulation*. Springer-Verlag, New York, 2007.
- [3] D. Austin. How google finds your needle in the web's haystack. *American Mathematical Society Feature Column*, 10(12), 2006. <http://www.ams.org/samplings/feature-column/fcarc-pagerank>.
- [4] K. Avrachenkov and N. Litvak. Decomposition of the Google PageRank and Optimal Linking Strategy. Rapport de recherche RR-5101, INRIA, 2004.
- [5] B. Benedelman and B. Lockwood. Measuring bias in "organic" web search. <http://www.benedelman.org/searchbias>, 2012.
- [6] D. P. Bertsekas and S. E. Shreve. *Stochastic Optimal Control: The Discrete Time Case*. Academic Press, New York, 1978.
- [7] F. BouSSION, P. Maillé, and B. Tuffin. Net neutrality debate: Impact of competition among ISPs. In *Proceedings of the Fourth International Conference on COMMunication Systems and NETWORKS (COMSNETS)*, Bangalore, India, 2012.
- [8] P. Coucheney, P. Maillé, and B. Tuffin. Comparison of search engines non-neutral and neutral behaviors. *SIGMETRICS Performance Evaluation Review*, 40(2):14–17, October 2012.
- [9] P. Coucheney, P. Maillé, and B. Tuffin. Network Neutrality Debate and ISP Inter-Relations: Traffic Exchange, Revenue Sharing, and Disconnection Threat. Rapport de recherche, Inria, January 2012.
- [10] J. Crowcroft. Net neutrality: The technical side of the debate: A white paper. *ACM SIGCOMM Computer Communication Review*, 7(1), January 2007.
- [11] R. Dejarnette. Click-through rate of top 10 search results in Google, 2012.
- [12] B. Edelman and B. Lockwood. Measuring bias in "organic" web search. <http://www.benedelman.org/searchbias/>, Jan 2011.

- [13] R. Frieden. Internet Packet Sniffing and its Impact on the Network Neutrality Debate and the Balance of Power between Intellectual Property Creators and Consumers. Technical report, SSRN, 2007. [http://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=995273](http://papers.ssrn.com/sol3/papers.cfm?abstract_id=995273).
- [14] W.A. Hanson and K. Kalyanam. *Internet Marketing and E-Commerce*. International student edition. Thomson/South-Western, Mason, Ohio, 2007.
- [15] Inria. Inria’s response to ARCEP consultation about network neutrality, 2012.
- [16] Auction Insights. Decoding eBay’s best match. <http://www.auctioninsights.info/decoding-ebays-best-match.html>, 2008.
- [17] T.M. Lenard and R.J. May (Eds.). *Net Neutrality or Net Neutering: Should Broadband Internet Services be Regulated*. Springer, 2006.
- [18] P. Maillé, E. Markakis, M. Naldi, G. Stamoulis, and B. Tuffin. Sponsored search auctions: an overview of research with emphasis on game theoretic aspects. *Electronic Commerce Research*, 12:265–300, 2012.
- [19] P. Maillé, P. Reichl, and B. Tuffin. Internet governance and economics of network neutrality. In Antonis Antoniadis and Burkhard Stiller, editors, *Telecommunications Economics - Selected Results of the COST Action IS605 Econtel*, pages 108–116. Lecture Notes in Computer Science 7216, Springer Verlag, 2012.
- [20] P. Maillé and B. Tuffin. *Telecommunication Network Economics: From Theory to Applications*. Cambridge University Press, 2014. to appear.
- [21] P. Njoroge, A. Ozdaglar, N. Stier-Moses, and G. Weintraub. Investment in two sided markets and the net neutrality debate. *Review of Network Economics*, 2014. Forthcoming.
- [22] A. Odlyzko. Network neutrality, search neutrality, and the never-ending conflict between efficiency and fairness in markets. *Review of Network Economics*, 8(1):40–60, 2009.
- [23] R. Pasupathy and S. Kim. The stochastic root finding problem: Overview, solutions, and open questions. *ACM Transactions on Modeling and Computer Simulation*, 21(3):Article 19, 2011.
- [24] R. Pollock. Is Google the next Microsoft: competition, welfare and regulation in online search. *Review of Network Economics*, 9(4):76–104, 2011.
- [25] Return On Now. 2013 search engine market share by country. <http://returnonnow.com/internet-marketing-resources/2013-search-engine-market-share-by-country>, 2013.
- [26] D. Rushe. Eric Schmidt Google senate hearing – as it happened. <http://www.guardian.co.uk/technology/blog/2011/sep/21/eric-schmidt-google-senate-hearing>, 2012.
- [27] H. Williams. Measuring search relevance. <http://www.ebaytechblog.com/2010/11/10/measuring-search-relevance>, 2010.
- [28] J. D. Wright. *Defining and Measuring Search Bias: Some Preliminary Evidence*. George Mason University School of Law, 2012.