

User Satisfaction in Competitive Sponsored Search

David Kempe
Department of Computer Science
University of Southern California
dkempe@usc.edu

Brendan Lucier
Microsoft Research
brlucier@microsoft.com

ABSTRACT

We present a model of competition between web search algorithms, and study the impact of such competition on user welfare. In our model, search providers compete for customers by strategically selecting which search results to display in response to user queries. Customers, in turn, have private preferences over search results and will tend to use search engines that are more likely to display pages satisfying their demands.

Our main question is whether competition between search engines increases the overall welfare of the users (i.e., the likelihood that a user finds a page of interest). When search engines derive utility only from customers to whom they show relevant results, we show that they differentiate their results, and every equilibrium of the resulting game achieves at least half of the welfare that could be obtained by a social planner. This bound also applies whenever the likelihood of selecting a given engine is a convex function of the probability that a user's demand will be satisfied, which includes natural Markovian models of user behavior.

On the other hand, when search engines derive utility from all customers (independent of search result relevance) and the customer demand functions are not convex, there are instances in which the (unique) equilibrium involves no differentiation between engines and a high degree of randomness in search results. This can degrade social welfare by a factor of $\Omega(\sqrt{N})$ relative to the social optimum, where N is the number of webpages. These bad equilibria persist even when search engines can extract only small (but non-zero) expected revenue from dissatisfied users, and much higher revenue from satisfied ones.

Categories and Subject Descriptors

H.4 [Information Systems Applications]: Miscellaneous

Keywords

search engine; competition; game theory; equilibrium; user satisfaction

1. INTRODUCTION

One of the most central tenets of economic market theory is that competition between vendors will benefit consumers, whether by driving down prices, encouraging product differentiation, or increasing the overall quality of goods. This belief is supported by many models and a long history of empirical studies. (See, e.g., [14, 20] for classic overviews.) In a market where consumers directly buy products from vendors, these benefits of competition are clear: given a choice between multiple competing products, a customer can be expected to choose whichever he finds most appealing at its given price. This dynamic naturally leaves inferior and overpriced products unsold.

In this paper, we analyze a game of competition between *search engines*, and its impact on the welfare of the users. The main difference between the market for sponsored search and classic oligopoly models is that in sponsored search, users are not the *customers* (since search results are provided to them for free), but the *product*, which is sold to advertisers. As a result, a search engine might have an interest in satisfying users only inasmuch as it can monetize their visit and satisfaction — this results in a potentially serious misalignment of incentives, and thus loss in user welfare.

At a high level, a competition between search engines is essentially a competition between (potentially randomized) *algorithms* that return ordered collections of webpages in response to search queries. A query typically does not completely determine which web pages are being searched for.¹ As a result, search engines will generally not be able to deterministically satisfy every user; in return, a searcher may not know whether a set of results will satisfy his demand before choosing a search provider.²

If there is only a single engine that acts as a monopolist, the algorithmic solution to the web search problem is clear: given a query, the engine should display the pages that maximize the probability that the user is satisfied. (We assume that engines have estimates of searchers' probabilities for desiring different pages.) When multiple engines compete for customers, it is less clear how the engines will behave at equilibrium. Will they differentiate their search results to capture different parts of the market, or will competition perversely drive them toward displaying esoteric webpages in their rankings, degrading the overall user experience?

The main insight of our work is that the answer depends on two features of the competition: (1) the alignment of

¹Think of polysemy, or simply different information needs of different users.

²In this sense, search results are an “experience good” [19].

utility between the engines and users, and (2) how users respond to search result quality. In order to elaborate on these insights, we briefly describe some features of our basic model of competitive search (The missing parts of the model are defined in Section 2.)

A brief overview of the model.

K search engines are competing for the attention of users, by displaying the set of N web pages in some order. Each query by a user³ induces a known probability distribution over subsets S of web pages; the probability associated with the set S is the probability that the user would be satisfied upon seeing at least one page in S . When the query arrives, each search engine can select (a distribution over) the order in which the N pages are displayed. In addition to the set S , the user's type also includes a patience threshold t : he is satisfied if the first t pages in the ordering contain at least one page from S .

Based on the distributions of orders in which the K search engines display pages, a user interested in a set S will choose an engine to visit. The choice can be probabilistic, and will be correlated in some way with search performance. We assume that the probability of choosing engine i is monotone non-decreasing in the probability that i would satisfy the user, and monotone non-increasing in the probabilities that the competing engines $i' \neq i$ would satisfy him. This model of user behavior is very general, allowing for many concrete instantiations of such *selection rules*, including natural Markovian models⁴, proportional choice, and others.

We note that even though a user's choice depends on the search engines' algorithms, the intended interpretation is not that users know their desired webpages in advance (which would eliminate the need for search engines). Rather, our model is intended to capture the fact that the fraction of users choosing an engine can depend, possibly in the long run, on the likelihood that a query on a certain topic will result in satisfactory search results.

It remains to describe the search engine payoffs in the game. When a user visits a search engine, it derives revenue by showing ads to the user. In sponsored search, the user is shown ads that are aligned with his perceived intent. Thus, even conditioned on a user having chosen a given engine, the revenue generated can be correlated with how well the engine inferred the user's intent, which is correlated with whether the user was satisfied by the search results. Motivated by this, we define engine utilities as follows: an engine receives a payoff of 1 for every user that visits the engine *and is satisfied*, and a payoff of $\beta \in [0, 1]$ for users that visit the engine but are not satisfied. That is, when $\beta = 1$ an engine receives full benefit from each user regardless of the search results, whereas $\beta = 0$ implies that an engine receives no revenue from a customer if it presents irrelevant results.

There is a natural interpretation of the parameter β in terms of the type of advertising chosen by the search engines. When using per-impression advertising (e.g., banner ads), the search engine receives the full payment from the advertiser when a user visits, regardless of satisfaction; this corresponds to $\beta = 1$. On the other hand, with per-click advertising, the search engine is paid only when a user clicks

on an ad, which will be highly correlated with the relevance of the results and ads displayed; thus, pure per-click advertising corresponds to $\beta = 0$. A mix of the two advertising models naturally gives rise to intermediate values of β .

We emphasize that even in the extreme case of $\beta = 1$ (a fully per-impression advertising model), user satisfaction is still relevant as it provides a way to attract users away from competing engines. The difference lies in how a user is valued, *given* that he has already made his choice about which engine to use.

Our Results.

We find that the outcome of competition depends heavily on both the value of β (i.e., on whether or not an engine derives direct benefit from unsatisfied users), and on the form of the selection rule (i.e., how users respond to changes in satisfaction probability). Our main analytical result, informally stated, is the following dichotomy:

1. If $\beta = 0$ (search engines obtain revenue only when a user is satisfied), or if the engine selection rule is convex, then the equilibrium (which is essentially unique) will have a large amount of specialization. Engines will choose deterministic strategies aimed at targeting particular subsets of the users. Regardless of the number of search engines, the Price of Anarchy is bounded by 2, even for mixed or correlated equilibria, and this bound is tight. In other words, competition degrades user satisfaction at most by a factor 2 compared to the outcome if engines were to pool their resources to best satisfy users.
2. In contrast, if $\beta > 0$, then for a large class of non-convex engine selection rules, there exist symmetric equilibria, i.e., all engines use the same algorithm. In these equilibria, there is no specialization, meaning that users are no better off than with a single engine. In fact, there are cases in which the equilibria can be *worse* than the optimal single-engine solution, with a pure Price of Anarchy of $\Omega(\sqrt{N})$ even when $K = 2$. When $\beta = 1$ we show that this equilibrium is unique, and hence even the pure Price of Stability is $\Omega(\sqrt{N})$.

Thus, informally stated, when search engines only want to satisfy viewers (e.g., because they use purely per-click advertising), searcher welfare is guaranteed to be close to the social optimum, whereas it can be very far from optimum when search engines only want to attract viewers (e.g., because they use purely per-impression advertising). However, these inefficiencies vanish under certain models of user behavior that result in convex selection rules. The intuition behind this phenomenon is that convexity incentivizes engines to have high satisfaction probability for a few user types, rather than infrequently satisfying many types. As a concrete example, we show that convexity occurs under a natural class of Markovian models of user behavior.

Related Work.

Our model is closely related to the study of competitive algorithm design due to Immorlica et al. [7]. In that work, two algorithms (for a general optimization problem) compete to maximize the probability that they will outperform their opponent on an (unknown, randomly selected) input. Competitive search is given as an example in [7], though

³whom we can imagine as drawn from an infinite population

⁴For example, a user might stick with his current engine until it fails to satisfy a query, at which point he switches to an alternative.

their model differs significantly from our own: in their version, the goal of each engine is to place a single desired page closer to the top of the search results than the opponent. The all-or-nothing nature of this goal creates some unintuitive best-response strategies, and may be questionable in terms of modeling actual web searcher behavior. In contrast, we focus on a richer model of competition between search algorithms, and explore the (in)efficiency of equilibria for this class of games.

A somewhat similar model of search engine competition was proposed by Khossainov and Kushmerick [9, 10]. They consider a metasearcher that directs queries to specialized engines, and propose and experimentally evaluate reinforcement learning approaches for the search engines.

Differentiation between competing search engines has been studied under economic models by Mukhodhyay, Telang, and Rajan [16]. They primarily focus on *vertical* differentiation, i.e., why lower-quality search engines can survive in a market, as well as a model of horizontal differentiation that somewhat resembles Hotelling’s model for competition [6].

Competition for user participation has been modeled in the context of auctions [13, 15, 17]. In these settings, each auctioneer selects a mechanism; bidders then select an auction to participate in and decide on their bidding strategy.

The revenue of a single search engine has been studied under various models of user behavior and advertising revenue [1, 3, 5, 8]. These works focus on the game played by advertisers, rather than by multiple search engines.

The analysis of horizontal product differentiation in competitive markets has a long history in economic theory. The classic Hotelling model [6] suggests that competing firms will tend to produce similar products. Later works, beginning with Aspremont, Gabszewicz, and Thisse [4], find that product differentiation can improve revenue by, for example, thinning markets to affect future prices. See [20] for a treatment of this classic theory.

Our analysis of competition in the case $\beta = 0$ shares some similarities with the analysis of valid utility games [21] and smooth games [18]. We note that the click-through variant of our game is mathematically similar to Oren and Kleinberg’s recent analysis of a research credit allocation game [11], where researchers select projects in which to invest effort. Besides the motivation behind these games, the primary difference is that we model a search engine’s (pure) strategy as a probability distribution over search results, rather than a deterministic list of pages.

2. MODELS AND CONCEPTS

We write $[n] = \{1, \dots, n\}$. Vectors are denoted by bold face. For a vector \mathbf{x} , we use \mathbf{x}_{-i} to denote the vector without coordinate i (extending the notation to sets of coordinates as well), and (y, \mathbf{x}_{-i}) for the vector with coordinate i replaced by y . For an event \mathcal{E} , we write $\mathbf{1}[\mathcal{E}]$ for the function that is 1 when \mathcal{E} happens and 0 otherwise.

Let \mathcal{A} be a finite set, and p a distribution over \mathcal{A} . We say that p is in *general position* if for every pair $S, S' \subseteq \mathcal{A}$ of distinct sets ($S \neq S'$), we have that $p(S) \neq p(S')$.

The Search Engine Competition Game.

The game is played by K engines. There is a universe Ω of N possible webpages, and an infinite population of *users*

(or *searchers*).⁵ There is a commonly known distribution Γ over user types, where $\gamma(S, t)$ is the probability of drawing the pair (S, t) . This distribution incorporates the search query issued by the user; in other words, we consider the game for a specific query, since games for different queries can be solved independently.⁶ Without loss of generality, we assume that $\sum_{S \ni n} \sum_t \gamma(S, t) > 0$ for all pages $n \in \Omega$, i.e., each page satisfies at least one type in the support of Γ .

The engines’ strategies are based on *permutations* of pages. When engine i plays permutation σ_i , she places page $\sigma_i(j)$ in slot j (for each j). We also think of σ_i as a *matching* between slots and pages, and write $\sigma_i[t] = \{\sigma_i(j) \mid j \in [t]\}$ for the set of pages placed in the first t slots by engine i . A user with type (S, t) is *satisfied* by a matching σ_i precisely if $\sigma_i[t] \cap S \neq \emptyset$, i.e., if at least one of the pages satisfying the user is among the top t pages displayed by the engine.

When a user visits engine i , the engine obtains utility 1 if the user is satisfied, and utility $\beta \in [0, 1]$, a fixed parameter of the game, if the user is not satisfied. The parameter β models the tradeoff between optimizing for market share vs. customer satisfaction, or the fraction of display ads among the ads shown by the engine.

A *pure strategy* for engine i is a distribution \mathcal{D}_i over matchings. (This apparent misnomer is discussed in more detail below.) A *(pure) strategy profile* $\mathcal{D} = (\mathcal{D}_1, \dots, \mathcal{D}_K)$ is a vector of strategies, one for each engine. A user type (S, t) and engine strategy \mathcal{D}_i together determine a satisfaction probability $q_i(S, t, \mathcal{D}_i) = \text{Prob}_{\sigma \sim \mathcal{D}_i}[\sigma[t] \cap S \neq \emptyset]$. We often write q_i for $q_i(S, t, \mathcal{D}_i)$ when the type and strategy are clear from the context. The vector of the satisfaction probabilities for all engines is denoted by $\mathbf{q}(S, t, \mathcal{D})$ or \mathbf{q} . Engines may also mix strategies, playing distributions over distributions.

Two permutations σ, σ' are *equivalent* with respect to Γ if, for every type (S, t) in the support of Γ , (S, t) is satisfied by either both or neither of the permutations. For example, if every user has a threshold of at least 2, then two permutations that differ only in the order of the first two pages are equivalent. A pure strategy is *deterministic* if its support consists entirely of permutations that are equivalent to each other; we abuse terminology by stating that such a strategy consists of a single permutation. A strategy profile (pure or mixed) is *deterministic* if the support of each engine’s distribution contains only deterministic strategies.

A user of type (S, t) must select which search engine to use, taking into account the profile \mathbf{q} of satisfaction probabilities. A *selection rule* is a function f which maps \mathbf{q} to a distribution over search engines. We write $f_i(\mathbf{q})$ for the probability that a user selects engine i given \mathbf{q} .

The precise sequence of events in the game is as follows. First, the engines jointly determine a mixed strategy profile \mathcal{C} . Let $\rho(\mathcal{D})$ be the density function of the joint distribution of pure strategies under \mathcal{C} . Notice that we allow the random choices to be correlated. Once \mathcal{C} is determined, a vector of pure (but not necessarily deterministic) strategies \mathcal{D} is drawn according to ρ . Next, a user type (S, t) is realized from Γ . \mathcal{D} is known to the user, and determines a satisfaction probability profile $\mathbf{q} = (q_1, \dots, q_K)$ for the engines. The user selects a search engine stochastically according to the distribution $f(\mathbf{q})$. After the user selects an engine, say i , a permutation σ_i is realized from the distribution \mathcal{D}_i . If

⁵We consistently use female pronouns for search engines and male pronouns for searchers.

⁶This holds under the (strong) assumption that users choose their engine per query. See a brief discussion of this issue in Section 6.

the user is satisfied by σ_i , engine i receives a payoff of 1; otherwise, she receives a payoff of β . All other engines $i' \neq i$ receive payoff 0. In summary, the payoff to engine i under the joint strategy profile \mathcal{C} is

$$\int_{\mathcal{D}} \sum_{\substack{S \subseteq \Omega \\ t \geq 0}} \gamma(S, t) \cdot f_i(\mathbf{q}(S, t, \mathcal{D})) \cdot (\beta + (1 - \beta)q_i(S, t, \mathcal{D})) \rho(\mathcal{D}) d\mathcal{D}. \quad (1)$$

We call this game the *Search Engine Competition Game* with parameters β and Γ , and denote it by $\text{SEARCH}_{\beta}(\Gamma)$.

Strategies, Equilibria, and the Price of Anarchy.

We now discuss in more depth the choice to consider distributions \mathcal{D}_i as *pure* strategies. The reason that we cannot simply consider engines as *mixing* over pure strategies which are *permutations* is that the engines' payoffs are defined in terms of user behavior, and users' behaviors depend on the entire probability distributions chosen by the engines.

In other words, it is necessary to know the distributions over pure strategies in order to even define the game's payoffs. For the same reason, when engines are playing mixed strategies, it is necessary to assume that the searchers are aware of the pure strategies drawn by the engines; this would be realistic if the changes in engines' strategies occurred at a much slower time scale than the users' adaptation.

We will carefully distinguish nomenclature between pure strategies and *deterministic* strategies, which — as defined above — are ones that have probability 1 of selecting a particular permutation (or its equivalence class). With this understanding in hand, a (*pure*) *Nash Equilibrium* is a strategy profile \mathcal{D} in pure strategies from which no single engine can unilaterally deviate to strictly improve her utility. Pure Nash Equilibria may not exist (even for $\beta = 0$) — we show a simple example of this in the full version of the paper ⁷ — so we broaden the space of allowed strategy profiles.

DEFINITION 1 (CORRELATED EQUILIBRIUM). *Fix a finite probability space (Ω, δ) . For each $i = 1, \dots, K$, let $Z_i : \Omega \rightarrow \mathcal{Z}_i$ be a signal (random variable), using the same domain, and write $\mathbf{Z} = (Z_1, \dots, Z_K)$.*

The strategy of player i is a mapping μ_i from \mathcal{Z}_i to the pure strategies of player i . $(\mu_1, \mu_2, \dots, \mu_K)$ is a correlated equilibrium iff for each player i and each mapping μ'_i from \mathcal{Z}_i to the pure strategies of player i , we have that

$$\begin{aligned} \mathbb{E}_{\mathbf{Z} \sim \delta} [u_i((\mu_1(Z_1), \dots, \mu_i(Z_i), \dots, \mu_K(Z_K)))] \\ \geq \mathbb{E}_{\mathbf{Z} \sim \delta} [u_i((\mu_1(Z_1), \dots, \mu'_i(Z_i), \dots, \mu_K(Z_K)))] . \end{aligned} \quad (2)$$

In other words, given the signal he received, each player prefers playing according to the suggested strategy over any other strategy.

We denote the set of all Correlated Equilibria by $\text{EQ}_{\beta}^{\text{Corr}}(\Gamma)$, or EQ^{Corr} when the parameters are clear from the context.

We measure the quality of a strategy profile by the welfare of the searchers, i.e., the probability that a user drawn from Γ is satisfied by his chosen engine. The welfare, $W(\mathcal{C})$, is

$$\int_{\mathcal{D}} \sum_{\substack{S \subseteq \Omega \\ t \geq 0}} \gamma(S, t) \cdot \sum_i f_i(\mathbf{q}(S, t, \mathcal{D})) \cdot q_i(S, t, \mathcal{D}) \cdot \rho(\mathcal{D}) d\mathcal{D}. \quad (3)$$

Let OPT denote a strategy profile for the engines that maximizes the searcher welfare. The degradation of welfare at equilibrium, compared to this optimum profile, is typically measured by two quantities: the (pure, mixed, or correlated) Price of Anarchy (PoA) [12, 18] $\sup_{\mathcal{D} \in \text{EQ}^{\text{Corr}}} \frac{W(\text{OPT})}{W(\mathcal{D})}$ captures the worst-case welfare loss resulting from selfish rational behavior by the engines; the Price of Stability (PoS) [2] $\inf_{\mathcal{D} \in \mathcal{N}} \frac{W(\text{OPT})}{W(\mathcal{D})}$ captures the welfare loss incurred by insisting that a strategy profile must be stable against selfish unilateral deviations. (Here, the infimum is usually taken over Nash Equilibria \mathcal{N} , as correlated equilibria are more permissive, and would thus lower the value.)

The Singleton Search Engine Game.

All of the lower bounds on PoA (or PoS) which we prove in this paper arise already in a very restricted form of the Search Engine Competition Game, which we call the *Singleton Search Engine Game*. In the Singleton Search Engine Game, each user always has a patience threshold of 1, and desires a singleton set. Thus, the distribution Γ can be specified simply as $\gamma(n) := \gamma(\{n\}, 1)$. Furthermore, since no user ever looks past the first slot, engine strategies are equivalent iff they agree in the page in the first slot. In particular, for the purpose of the game, a distribution \mathcal{D} is completely specified by the probability distribution of pages in slot 1. We will therefore frequently speak simply of an engine “displaying” or “choosing” a page, without specifying a slot (which is understood to be slot 1).

This distribution is characterized by the probability of satisfying any given type of user (which can now be associated with the desired page n). Therefore, for instances of the Singleton Search Engine Game, an engine i 's pure strategy can be fully specified by the probability vector $\mathbf{q}_i = (q_i(n))_{n=1}^N$ of probabilities of displaying pages $n = 1, \dots, N$.

User Behavior and Selection Rules.

We model the users as stochastic entities, choosing engines via a probabilistic process that depends on their probability of being satisfied. A selection rule f captures a model of user behavior. We assume that the same selection rule is used by each user,⁸ and that the effect of the engines' strategies on user behavior can be captured entirely by the satisfaction probability profile \mathbf{q} . We assume that each f_i is monotone non-decreasing in q_i and non-increasing in $q_{i'}$ for each $i' \neq i$. We do not assume that f_i is differentiable or continuous.

For some of our results, we will be interested in selection rules f that are *symmetric* in the following sense: the value of $f_i(\mathbf{q})$ is invariant under relabeling of the engines. When analyzing symmetric rules, we write $f(\mathbf{q}) = f_1(\mathbf{q})$.

This model is very general, encompassing almost all rational selection rules one could imagine. We illustrate it with several examples, capturing different models of behavior.

Proportional selection rule: $f_i(\mathbf{q}) = \frac{q_i}{\sum_j q_j}$, with $f_i(\mathbf{0}) := 1/K$. In this rule, a user picks engine i with probability proportional to the probability that he will be satisfied by engine i . Similar rules have frequently been used in past

⁷available on the arXiv as 1310.4098

⁸We note that the space of selection rules is convex, so the selection rule in which each user independently randomizes between using a rule f_1 or an alternative rule f_2 is, itself, a selection rule.

models that involved a user selecting a vendor according to the quality of the vendors' products.

Markovian selection rule: $f_i(\mathbf{q}) = \frac{b_i}{\sum_j b_j}$, where $b_j = \frac{1}{1-q_j}$ for each engine j . If any q_j is equal to 1, then $f_i(\mathbf{q}) = 1/k$ if $q_i = 1$, where k is the number of engines with $q_j = 1$, and $f_i(\mathbf{q}) = 0$ if $q_i < 1$. This selection rule captures the stationary distribution of a process in which a user repeatedly uses a given engine until she does not satisfy him, then switches to another engine uniformly at random. (A more general version is defined in Section 5.)

Majority selection rule: Let $A = \operatorname{argmax}_i q_i$ be the set of engines maximizing the satisfaction probability of a given user. Then the Majority selection rule sets $f_i(\mathbf{q}) = 1/|A|$ for all $i \in A$, and $f_i(\mathbf{q}) = 0$ for all $i \notin A$. The Majority selection rule captures the behavior of a user who chooses an engine that maximizes his probability of being satisfied, breaking ties uniformly at random.

3. CONVEXITY, DETERMINISM, AND HIGH WELFARE

In this section, we examine the role of convexity in the search engines' utility functions. Roughly speaking, we show that when the selection rule f is convex, each engine's utility is strictly convex. This convexity, or alternatively, the assumption that $\beta = 0$, implies that every correlated equilibrium is deterministic; in other words, engines do not mix over distributions of permutations, but only over permutations. For pure equilibria, this implies that engines simply play one permutation deterministically. In turn, we will show that the restriction to deterministic strategies implies a bound of 2 on the correlated PoA.

In Section 5, we will show that, for a natural and general class of Markovian user behaviors, the selection rule f is indeed convex; hence, for arbitrary β , this type of user behavior yields a correlated PoA of at most 2.

We begin by defining a non-indifference property of selection rules. Roughly speaking, we wish to exclude selection rules for which a search engine's probability of satisfying a user has effectively no impact on that user's selection. We define the *support* $S(\mathbf{q})$ of a satisfaction probability vector \mathbf{q} to be the set of engines i for which $q_i > 0$. We say that f is *non-indifferent* if the following two properties hold. First, whenever $|S(\mathbf{q})| \geq 2$ and $q_j < 1$ for some $j \in S(\mathbf{q})$, there exists some $i \in S(\mathbf{q})$ such that $f_i(\mathbf{q}) < f_i(1, \mathbf{q}_{-i})$. That is, if at least two engines would satisfy a user with positive probability, and not all supporting engines satisfy the user with probability 1, then at least one supporting engine would be more likely to attract the user by sufficiently increasing her probability of satisfying him. Second, if there exists some i such that $q_i = 1$, then $f_j(\mathbf{q}) = 0$ for all $j \notin S(\mathbf{q})$. That is, if some engine is guaranteed to satisfy the user, then the user will not choose an engine that is guaranteed *not* to satisfy him. The non-indifference property is quite weak; in particular, it is significantly weaker than assuming that f_i is strictly increasing in q_i . For example, both the Majority and Markovian selection rules described at the end of Section 2 are not strictly increasing in the satisfaction probabilities, but do satisfy non-indifference.

THEOREM 3.1. *Suppose that Γ is in general position, and at least one of the following three holds for each i :*

1. $\beta < 1$, and $f_i(q_i, \mathbf{q}_{-i})$ is a convex function of q_i ;
2. $\beta = 1$, and $f_i(q_i, \mathbf{q}_{-i})$ is strictly convex in q_i ;
3. $\beta = 0$ and f_i is non-indifferent.

Then, each correlated equilibrium of the game $\text{SEARCH}_\beta(\Gamma)$ is deterministic.

The proof of Theorem 3.1 (in the appendix) follows from the observation that if selection rules are convex or $\beta = 0$, an engine can increase her utility by choosing two outcomes she randomizes between, deciding which one contributes more to the overall utility, and increasing the probability of returning that outcome at the expense of the other.

We now show that if engines apply deterministic strategies at equilibrium, the resulting outcomes will generate high social welfare. Combining this result with Theorem 3.1, we conclude that the PoA of $\text{SEARCH}_\beta(\Gamma)$ is at most 2 under the conditions of Theorem 3.1.

THEOREM 3.2. *Suppose that f is non-indifferent and every correlated equilibrium of $\text{SEARCH}_\beta(\Gamma)$ is deterministic. Then, the correlated PoA of $\text{SEARCH}_\beta(\Gamma)$ is at most 2.*

Proof. Consider a correlated deterministic equilibrium of $\text{SEARCH}_\beta(\Gamma)$, characterized by a joint distribution $\rho(\sigma)$ over vectors of permutations σ . Let $\hat{\sigma}$ be the vector of permutations for each engine under the socially optimal profile. Focus on one engine i and one signal Z_i she receives. Let σ_i be the permutation played by i when seeing Z_i . From the perspective of i , conditioned on receiving Z_i , the other engines are playing a correlated distribution $\rho_{-i}(\sigma_{-i})$.

Let $\mu_{S,t}^{Z_i} = \text{Prob}[\forall i' : \sigma_{i'}[t] \cap S = \emptyset \mid Z_i]$ be the probability that no engine i' (including i) satisfies the searcher type (S, t) , given that i received the signal Z_i . Consider a deviation by i to any $\hat{\sigma}_j$. This will certainly attract any searcher who would not be satisfied by any other engine, but would be satisfied by $\hat{\sigma}_j$. Also, any other searcher satisfied by $\hat{\sigma}_j$ would select engine i with probability at least $f_i(1)$, by the assumed monotonicity of the selection rule. Finally, this deviation will not affect the behavior of searchers that are not satisfied by any engine or $\hat{\sigma}_j$. Putting this together, the expected conditional utility of this deviation is at least

$$\begin{aligned} \mu_j = & \sum_{S \subseteq \Omega} \sum_{t \geq 0} \gamma(S, t) \cdot \mathbf{1}[\hat{\sigma}_j[t] \cap S \neq \emptyset] \cdot \mu_{S,t}^{Z_i} \\ & + \sum_{S \subseteq \Omega} \sum_{t \geq 0} \gamma(S, t) \cdot f_i(1) \cdot \mathbf{1}[\hat{\sigma}_j[t] \cap S \neq \emptyset] \cdot (1 - \mu_{S,t}^{Z_i}) \\ & + \beta \cdot \sum_{S \subseteq \Omega} \sum_{t \geq 0} \gamma(S, t) \cdot f_i(0) \cdot \mathbf{1}[\hat{\sigma}_j[t] \cap S = \emptyset] \cdot \mu_{S,t}^{Z_i}. \end{aligned}$$

Since engine i , receiving the signal Z_i , does not want to deviate to $\hat{\sigma}_j$ at the correlated equilibrium, μ_j is at most

$$\begin{aligned} & \int_{\sigma_{-i}} \sum_{S \subseteq \Omega} \sum_{t \geq 0} \gamma(S, t) \cdot f_i(\mathbf{q}(S, t, \sigma)) \cdot q_i(S, t, \sigma) \cdot \rho_{-i}(\sigma_{-i}) d\sigma_{-i} \\ & + \beta \cdot \sum_{S \subseteq \Omega} \sum_{t \geq 0} \gamma(S, t) \cdot f_i(0) \cdot \mu_{S,t}^{Z_i}. \end{aligned}$$

The first term is the utility that i obtains from searchers she attracts and satisfies, while the second term is the utility obtained from searchers that are not satisfied by any

engine, and visit engine i . Taking an expectation over signals Z_i , writing $\mu_{S,t} = \mathbb{E}_{Z_i \sim \rho_i} [\mu_{S,t}^{Z_i}]$, and rearranging the terms with β factors, we conclude

$$\begin{aligned} & \sum_{S \subseteq \Omega} \sum_{t \geq 0} \gamma(S, t) \cdot \mathbb{1}[\hat{\sigma}_j[t] \cap S \neq \emptyset] \cdot \mu_{S,t} \\ & + \sum_{S \subseteq \Omega} \sum_{t \geq 0} \gamma(S, t) \cdot f_i(\mathbf{1}) \cdot \mathbb{1}[\hat{\sigma}_j[t] \cap S \neq \emptyset] \cdot (1 - \mu_{S,t}) \\ & \leq \int_{\sigma} \sum_{S \subseteq \Omega} \sum_{t \geq 0} \gamma(S, t) \cdot f_i(\mathbf{q}(S, t, \sigma)) \cdot q_i(S, t, \sigma) \cdot \rho(\sigma) d\sigma \\ & + \beta \cdot \sum_{S \subseteq \Omega} \sum_{t \geq 0} \gamma(S, t) \cdot f_i(\mathbf{0}) \cdot \mathbb{1}[\hat{\sigma}_j[t] \cap S \neq \emptyset] \cdot \mu_{S,t}. \end{aligned}$$

Taking a sum over all i , and applying the definition of equilibrium welfare $W(\rho)$, we get

$$\begin{aligned} & K \cdot \sum_{S \subseteq \Omega} \sum_{t \geq 0} \gamma(S, t) \cdot \mathbb{1}[\hat{\sigma}_j[t] \cap S \neq \emptyset] \cdot \mu_{S,t} \\ & + \sum_{S \subseteq \Omega} \sum_{t \geq 0} \gamma(S, t) \cdot \mathbb{1}[\hat{\sigma}_j[t] \cap S \neq \emptyset] \cdot (1 - \mu_{S,t}) \cdot \sum_i f_i(\mathbf{1}) \\ & \leq W(\rho) \\ & + \beta \cdot \sum_{S \subseteq \Omega} \sum_{t \geq 0} \gamma(S, t) \cdot \mu_{S,t} \cdot \mathbb{1}[\hat{\sigma}_j[t] \cap S \neq \emptyset] \cdot \sum_i f_i(\mathbf{0}). \end{aligned}$$

Using $\sum_i f_i(\mathbf{0}) \leq 1$, $\sum_i f_i(\mathbf{1}) = 1$, and summing over all j ,

$$\begin{aligned} & K \cdot W(\rho) \\ & \geq (K - \beta) \cdot \sum_{S \subseteq \Omega} \sum_{t \geq 0} \gamma(S, t) \mu_{S,t} \sum_j \mathbb{1}[\hat{\sigma}_j[t] \cap S \neq \emptyset] \\ & + \sum_{S \subseteq \Omega} \sum_{t \geq 0} \gamma(S, t) (1 - \mu_{S,t}) \sum_j \mathbb{1}[\hat{\sigma}_j[t] \cap S \neq \emptyset] \\ & \geq (K - \beta) \cdot \sum_{S \subseteq \Omega} \sum_{t \geq 0} \gamma(S, t) \mu_{S,t} \cdot \mathbb{1}[\exists j: \hat{\sigma}_j[t] \cap S \neq \emptyset] \\ & + \sum_{S \subseteq \Omega} \sum_{t \geq 0} \gamma(S, t) (1 - \mu_{S,t}) \cdot \mathbb{1}[\exists j: \hat{\sigma}_j[t] \cap S \neq \emptyset] \\ & = (K - \beta) (W(\text{OPT}) - \lambda \cdot W(\rho)) + \lambda \cdot W(\rho) \end{aligned}$$

for some $\lambda \in [0, 1]$. The last equality follows because

$$\begin{aligned} & \sum_{S \subseteq \Omega} \sum_{t \geq 0} \gamma(S, t) \cdot \mathbb{1}[\exists j: \hat{\sigma}_j[t] \cap S \neq \emptyset] = W(\text{OPT}) \quad \text{and} \\ & \sum_{S \subseteq \Omega} \sum_{t \geq 0} \gamma(S, t) (1 - \mu_{S,t}) \leq W(\rho). \end{aligned}$$

Rearranging, we have $W(\rho) \geq \frac{K-\beta}{(1+\lambda)K-\lambda(\beta+1)} \cdot W(\text{OPT})$, with $\lambda \in [0, 1]$. As $K \geq 2$ and $\beta \in [0, 1]$, this expression is minimized at $\lambda = 1$; hence, $W(\rho) \geq \frac{K-\beta}{2K-\beta-1} W(\text{OPT}) \geq \frac{1}{2} W(\text{OPT})$, as claimed. ■

The bound of 2 on the PoA is asymptotically tight, even for the pure PoA of the Singleton Search Engine Game with a convex selection rule and $\beta = 0$. In the tight instance, there are K pages, and $\gamma(1) = \frac{K+1}{2K}$, while $\gamma(n) = \frac{1}{2K}$ for all $n > 1$. Consider the profile in which each engine displays page 1. For any symmetric selection rule (convex or otherwise), each engine has utility $\frac{K+1}{K} \cdot \frac{1}{2K}$, which exceeds the utility $\frac{1}{2K}$ from switching to another page n . Hence, the strategy profile is an equilibrium (in fact, the unique one). It satisfies a $\frac{K+1}{2K}$ fraction of users. On the other hand, the

profile in which engine i displays page i satisfies all users. As K becomes large, the ratio approaches 2.

4. SYMMETRIC EQUILIBRIA AND LOW WELFARE

Next, we analyze equilibria of the Search Engine Competition Game when the conditions of Theorem 3.1 do not apply; that is, when $\beta > 0$ and the selection rule is not convex. We will show that the resulting misalignment of incentives can lead to equilibria with very poor social welfare, i.e., high PoA. All of our results are proved already for the Singleton Search Engine Game⁹, and all proofs are deferred to the appendix or the full version.

4.1 When Satisfaction is Irrelevant

We begin by considering the case $\beta = 1$. Our first result is that when the selection rule f is symmetric and satisfies a cross-concavity assumption, every pure Nash Equilibrium of the Singleton Search Engine Game is *symmetric*. In other words, search engines do not specialize: all engines apply the same strategy. A user can thus do no better than selecting an engine arbitrarily, and hence competition does not positively impact social welfare.

A selection rule f is *strictly cross-concave* if it satisfies the following property: for all x and y such that $x < y$, there exists a $\delta > 0$ such that for all $0 < \epsilon < \delta$,

$$\begin{aligned} & f((x + \epsilon, y, \mathbf{q}_{-\{1,2\}})) - f((x, y, \mathbf{q}_{-\{1,2\}})) \\ & > f((x, y, \mathbf{q}_{-\{1,2\}})) - f((x, y - \epsilon, \mathbf{q}_{-\{1,2\}})). \end{aligned}$$

Intuitively, if f is cross-concave, the increase in user visits that a (trailing) engine experiences when increasing her satisfaction probability is at least as large as the loss in visits incurred by a (leading) engine from decreasing her satisfaction probability by the same amount. That is, catching up in user visits is easier than expanding a lead.

THEOREM 4.1. *Let \mathbf{Q} be a pure strategy equilibrium of the Singleton Search Engine Game with $\beta = 1$ and symmetric and strictly cross-concave selection function f . Then \mathbf{Q} must be symmetric.*

The proof proceeds by directly arguing that if two engines employ different strategies, then cross-concavity implies that one of them can strictly improve her welfare by shifting her page distribution slightly toward the strategy used by the other. Theorem 4.1 implies that for symmetric and strictly cross-concave selection rules and $\beta = 1$, more engines do not provide any advantage in pure strategy equilibria. As we show below, this implies a PoA of $\Omega(K)$. However, we will show that the effects of competition can be even more negative, by giving lower bounds on the *Price of Stability* for the proportional selection rule. Thereto, we next characterize the (unique) equilibrium for the proportional selection rule. The proof follows from an analysis of the equilibrium condition that no engine wishes to shift probability mass from any page to any other.

THEOREM 4.2. *When $\beta = 1$ and f is the proportional selection rule, the Singleton Search Engine Game has a unique*

⁹The strong lower bounds continue to hold under the (realistic) assumption that $t \ll N$. When $t \approx N$, searchers will typically see most pages, so good welfare is obtained.

pure Nash Equilibrium where each engine selects page $n \in \Omega$ with probability $\gamma(n)$.

We use Theorem 4.2 to bound the PoS for the Search Engine Competition Game.

PROPOSITION 4.3. *Consider the Singleton Search Engine Game with $\beta = 1$ and the proportional selection rule.*

1. *There are instances in which the PoS is at least K .*
2. *There are instances with only two engines in which the PoS is $\Omega(\sqrt{N})$.*

The idea behind the second example of Proposition 4.3 is a type distribution in which two “popular” pages are each desired with probability $\Theta(1/\sqrt{N})$, while the remaining pages have probability $\Theta(1/N)$. The unique equilibrium then selects each page proportional to its probability, whereas the optimal outcome would have each engine displaying one of the popular pages deterministically.

4.2 When Satisfaction is Partially Beneficial

We now consider the intermediate case, $\beta \in (0, 1)$, which represents scenarios in which search engines derive some, but not full, expected advertising benefit from a user who is not served a desired page. We show that, for any fixed β , our lower bounds on the PoS in Proposition 4.3 can persist as long as the number of pages N is large. (However, for intermediate β values, we cannot rule out the existence of other equilibria, so we obtain lower bounds on the PoA, not the PoS.) Our conclusion is that the constant PoA bounds at $\beta = 0$ in general do not survive perturbations of β when the selection rule is not convex.

We again focus on the Singleton Search Engine Game with the proportional selection rule f . We first show that the existence of symmetric equilibria for this class, established for $\beta = 1$ in Section 4.1, continues to hold under small perturbations of β . The proof is conceptually similar to that of Theorem 4.2.

THEOREM 4.4. *Suppose that $\beta \in (0, 1)$. The Singleton Search Engine Game with proportional selection rule f has at most one symmetric equilibrium. Moreover, whenever $\beta > 1 - 1/K$, a symmetric equilibrium exists.*

In fact, for specific instances, symmetric equilibria (and thus high PoA) arise for much smaller values of β already. Specifically, we show that the bad examples from Proposition 4.3 persist for $\beta = \Omega(K/N)$ and $\beta = \Omega(K/\sqrt{N})$, respectively. Notice that since $N \gg K$, this means that in realistic settings, we cannot rule out high inefficiency due to search engine competition for any constant value of β . We emphasize that while Proposition 4.3 provided lower bounds on the *Price of Stability*, we have not ruled out the existence of asymmetric equilibria for $\beta \in (0, 1)$; therefore, the propositions in this section only prove lower bounds on the *Price of Anarchy*.

PROPOSITION 4.5. *Consider the Singleton Search Engine Game with proportional selection rule f . If $\beta > \frac{2K}{N+2K}$, there are instances with PoA $\Omega(K)$.*

PROPOSITION 4.6. *Consider the Singleton Search Engine Game with proportional selection rule f . If $\beta > \frac{6}{\sqrt{N}}$, there are instances with $K = 2$ and PoA $\Omega(\sqrt{N})$.*

5. CONVEXITY OF MARKOV RULES

We now describe a natural model of user behavior under which the selection rule is convex. By Theorems 3.1 and 3.2, if users behave according to this model, searcher welfare is high. The model generalizes the Markov selection rule described in Section 2. Each search engine i corresponds to a unique state. The user performs a walk on states, and always uses the engine corresponding to his current state. The transition probabilities between states depend on whether the user was satisfied in the previous round; we use $\tau_{i,j}^{(s)}$ and $\tau_{i,j}^{(f)}$ to denote the probability of moving from state i to state j upon satisfaction or failure (non-satisfaction), respectively. The process has a stationary distribution π , and the induced Markovian selection rule is $f_i(\mathbf{q}) = \pi_i$.

Leaving state i after a successful (resp., failed) query induces a distribution over the process’s state in the next round. That state defines an expected return time $R_i^{(s)}$ (resp., $R_i^{(f)}$) to state i : the expected number of rounds before the process returns to state i . Note that $R_i^{(s)}, R_i^{(f)} \geq 1$.

We call a Markov process *monotone* if, for all i , $\tau_{i,i}^{(s)} \geq \tau_{i,i}^{(f)}$ and $R_i^{(s)} \leq R_i^{(f)}$. The first condition states that a user is no more likely to leave his current engine after a successful query than after a failed query. The second condition says that, conditioning on having left engine i , the last query having been successful should not increase the expected return time. The Markov process is *strictly monotone* if the first inequality is strict. We show that monotone Markovian selection rules are convex.

THEOREM 5.1. *Any (strictly) monotone Markovian selection rule is (strictly) convex.*

Proof. For a given Markov process, consider the time spent in state i . Write $p_i^{(f)} = 1 - \tau_{i,i}^{(f)}$ and $p_i^{(s)} = 1 - \tau_{i,i}^{(s)}$ for the probability that the process leaves state i following a failure or success, respectively. Given that the process is in state i , the expected number of rounds until the process leaves state i is $\frac{1}{(1-q_i)p_i^{(f)} + q_i p_i^{(s)}}$. Conditioned on having left state i , the probability that the last search was successful is $\frac{q_i p_i^{(s)}}{(1-q_i)p_i^{(f)} + q_i p_i^{(s)}}$, and the probability that it failed is $\frac{(1-q_i)p_i^{(f)}}{(1-q_i)p_i^{(f)} + q_i p_i^{(s)}}$. We conclude that the expected return time, conditioned on having left state i , is

$$\frac{q_i p_i^{(s)}}{(1-q_i)p_i^{(f)} + q_i p_i^{(s)}} R_i^{(s)} + \frac{(1-q_i)p_i^{(f)}}{(1-q_i)p_i^{(f)} + q_i p_i^{(s)}} R_i^{(f)}.$$

Since the fraction of time spent in state i is the expected stay length divided by expected stay length plus return time, we have that π_i is equal to

$$\frac{\frac{1}{(1-q_i)p_i^{(f)} + q_i p_i^{(s)}}}{\frac{1}{(1-q_i)p_i^{(f)} + q_i p_i^{(s)}} + \frac{q_i p_i^{(s)} R_i^{(s)}}{(1-q_i)p_i^{(f)} + q_i p_i^{(s)}} + \frac{(1-q_i)p_i^{(f)} R_i^{(f)}}{(1-q_i)p_i^{(f)} + q_i p_i^{(s)}}}$$

We must show that π_i is a valid selection rule, meaning that it is non-decreasing in q_i . Taking the derivative with respect to q_i , we have

$$\frac{d}{dq_i} \pi_i = \frac{p_i^{(f)} R_i^{(f)} - p_i^{(s)} R_i^{(s)}}{(1 + p_i^{(f)} R_i^{(f)} (1 - q_i) + p_i^{(s)} R_i^{(s)} q_i)^2},$$

which is non-negative since $p_i^{(f)} \geq p_i^{(s)}$ and $R_i^{(f)} \geq R_i^{(s)}$. Next, we show that π_i is a convex function of q_i . Taking the second derivative with respect to q_i , we have

$$\frac{d^2}{dq_i^2} \pi_i = \frac{2(p_i^{(f)} R_i^{(f)} - p_i^{(s)} R_i^{(s)})^2}{(1 + p_i^{(f)} R_i^{(f)} (1 - q_i) + p_i^{(s)} R_i^{(s)} q_i)^3},$$

which is non-negative, and strictly positive if $p_i^{(f)} > p_i^{(s)}$. ■

6. CONCLUSIONS AND FUTURE WORK

We introduced a natural model of competition between multiple search algorithms, who are vying for visits from users searching for web pages. Our model interpolates between different objectives for the search engines. At one extreme is the objective of attracting as many searchers as possible. At the other extreme is an objective of attracting *and satisfying* as many searchers as possible.

Our main result is a strong dichotomy: when engines are directly incentivized to satisfy users, we expect to see specialization, leading to a market that serves users well and achieves a (tight) Price of Anarchy of 2. Otherwise, search engines will generally play symmetrically, duplicating each other's strategies. The Price of Stability in this setting can be as bad as $\Omega(K)$ or $\Omega(\sqrt{N})$. However, these pessimistic examples vanish under natural models of user behavior in which the demand for an engine is a convex function of that engine's search result quality.

A number of questions remain open. Our lower bounds apply to the pure Price of Stability; we did not rule out the existence of more benign mixed (or correlated) equilibria. Can the lower bounds be generalized? Our understanding of equilibria for $\beta \in (0, 1)$ is even more limited; how can we characterize the equilibria in this case, and can we rule out asymmetric pure equilibria?

One could also consider approximate equilibria, in which engines cannot unilaterally deviate to improve their payoffs *much*. It would be interesting to study the Price of Anarchy/Stability for approximate equilibria.

Finally, another layer of complexity may be added to the model to be more realistic. Our model posited that a user would choose an engine i for each search independently. In reality, most users are loyal to one search engine. This could be modeled by assuming that each user has a distribution over potential queries, and picks one engine to maximize the expected probability of satisfaction across all queries. Analyzing the Price of Anarchy/Stability of this more complex game is an interesting direction for future work.

Acknowledgments

We thank Shaddin Dughmi, Matt Elliott, Xinran He and Nicole Immorlica for useful discussions, and anonymous reviewers for helpful feedback. Work done in part while David Kempe was visiting Microsoft Research, New England.

APPENDIX

A. OMITTED PROOFS

Proof of Theorem 3.1: Fix a correlated equilibrium \mathcal{C} , and assume for contradiction that the support of \mathcal{C} is not entirely deterministic. Fix an engine i and a signal Z_i given to i with positive probability, such that conditioned on receiving Z_i , engine i plays a non-deterministic strategy

\mathcal{D}_i . From engine i 's viewpoint, conditioned on Z_i , the other engines $i' \neq i$ are playing mixed strategies, randomizing in possibly correlated ways over possibly non-deterministic strategies. Let $\rho_{-i}(\mathcal{D}_{-i})$ denote the joint density over the strategies played by engines $i' \neq i$, conditioned on Z_i .

Let $\text{Supp}(\Gamma)$ denote the support of Γ , and consider a type $(S, t) \in \text{Supp}(\Gamma)$. When engine i satisfies users of type (S, t) with probability x , [her](#) expected utility conditioned on the user type being (S, t) , $g_i(x; S, t)$, is

$$\int_{\mathcal{D}_{-i}} f_i((x, \mathbf{q}_{-i}(S, t, \mathcal{D}_{-i}))) \cdot (\beta + (1 - \beta) \cdot x) \cdot \rho_{-i}(\mathcal{D}_{-i}) d\mathcal{D}_{-i}.$$

If $f_i((x, \mathbf{q}_{-i}(S, t, \mathcal{D}_{-i})))$ is a convex function of x and $\beta < 1$, or if f_i is strictly convex in x , then for each (S, t) , $f_i((x, \mathbf{q}_{-i}(S, t, \mathcal{D}_{-i}))) \cdot (\beta + (1 - \beta) \cdot x)$ is a strictly convex function of x , and hence so is $g_i(x; S, t)$. Therefore, under either of the first two assumptions of the theorem, $g_i(\cdot; S, t)$ is strictly convex. (The analysis in the third case $\beta = 0$ will proceed similarly, but not explicitly be based on strict convexity.)

The utility of engine i from playing at this equilibrium is

$$u_i(\mathcal{D}) = \sum_{S \subseteq \Omega} \sum_{t \geq 0} \gamma(S, t) \cdot g_i(q_i(S, t, \mathcal{D}_i); S, t).$$

Engine i plays the non-deterministic strategy \mathcal{D}_i . For each permutation σ , let $\mathcal{D}_i(\sigma)$ be the probability that engine i chooses the permutation σ . Recall from Section 2 that

$$q_i(S, t, \mathcal{D}_i) = \sum_{\sigma} \mathcal{D}_i(\sigma) \cdot \mathbb{1}[\sigma[t] \cap S \neq \emptyset]; \quad (4)$$

thus, $q_i(S, t, \mathcal{D}_i)$ is a convex combination of the $\mathbb{1}[\sigma[t] \cap S \neq \emptyset]$ terms. In the first two cases of the theorem, we can use the convexity of g_i to obtain that

$$g_i(q_i(S, t, \mathcal{D}_i); S, t) \leq \sum_{\sigma} \mathcal{D}_i(\sigma) \cdot g_i(\mathbb{1}[\sigma[t] \cap S \neq \emptyset]; S, t). \quad (5)$$

In the third case ($\beta = 0$, no assumptions on f), Inequality (5) holds for a slightly different reason: we can write

$$\begin{aligned} & g_i(q_i(S, t, \mathcal{D}_i); S, t) \\ &= \int_{\mathcal{D}_{-i}} f_i((q_i(S, t, \mathcal{D}_i), \mathbf{q}_{-i}(S, t, \mathcal{D}_{-i}))) \\ & \quad \cdot q_i(S, t, \mathcal{D}_i) \cdot \rho_{-i}(\mathcal{D}_{-i}) d\mathcal{D}_{-i} \\ &\leq \int_{\mathcal{D}_{-i}} f_i((1, \mathbf{q}_{-i}(S, t, \mathcal{D}_{-i}))) \\ & \quad \cdot \sum_{\sigma} \mathcal{D}_i(\sigma) \cdot \mathbb{1}[\sigma[t] \cap S \neq \emptyset] \cdot \rho_{-i}(\mathcal{D}_{-i}) d\mathcal{D}_{-i} \\ &= \sum_{\sigma} \mathcal{D}_i(\sigma) \cdot \int_{\mathcal{D}_{-i}} f_i((\mathbb{1}[\sigma[t] \cap S \neq \emptyset], \mathbf{q}_{-i}(S, t, \mathcal{D}_{-i}))) \\ & \quad \cdot \mathbb{1}[\sigma[t] \cap S \neq \emptyset] \cdot \rho_{-i}(\mathcal{D}_{-i}) d\mathcal{D}_{-i} \\ &= \sum_{\sigma} \mathcal{D}_i(\sigma) \cdot g_i(\mathbb{1}[\sigma[t] \cap S \neq \emptyset]; S, t). \end{aligned}$$

For the inequality, we used the selection rule's monotonicity. In the penultimate step, we used that $\mathbb{1}[\sigma[t] \cap S \neq \emptyset] = 0$ whenever $\mathbb{1}[\sigma[t] \cap S \neq \emptyset] \neq 1$, so the whole expression for the particular σ is 0 anyway.

Next, we argue that Inequality (5) is in fact strict for at least one pair (S, t) with $\gamma(S, t) > 0$, in all cases. Because \mathcal{D}_i is not a deterministic strategy, there exist two

non-equivalent permutations σ, σ' with $\mathcal{D}_i(\sigma) \in (0, 1)$ and $\mathcal{D}_i(\sigma') \in (0, 1)$. Let $X \subseteq \text{Supp}(\Gamma)$ be the set of types satisfied by σ but not by σ' , and $X' \subseteq \text{Supp}(\Gamma)$ the set of types satisfied by σ' but not by σ ; note that $X \cup X' \neq \emptyset$ by non-equivalence. We then have $0 < q_i(S, t, \mathcal{D}_i) < 1$ for all $(S, t) \in X \cup X'$. Therefore, in Equation (4) for those types (S, t) , there are both terms with $\mathbb{1}[\sigma[t] \cap S \neq \emptyset] = 0$ and with $\mathbb{1}[\sigma[t] \cap S \neq \emptyset] = 1$. Thus, for the first two cases (when $g_i(\cdot; S, t)$ is strictly convex), Inequality (5) is strict for this (S, t) pair.¹⁰

Again, the reason for the strictness is slightly different in the third case. Focus on the searcher types $(S, t) \in X \cup X'$. Assume for contradiction that $q_j(S, t, \mathcal{D}_j) = 0$ for all $j \neq i$ and $(S, t) \in X \cup X'$, and that Inequality (5) is not strict. Thus, for each $(S, t), (S', t') \in X \cup X'$,

$$f_i(\mathbf{q}(S, t, \mathcal{D})) = f_i(1, \mathbf{0}) = f_i(\mathbf{q}(S', t', \mathcal{D})).$$

Because Γ is in general position, we know $\sum_{(S, t) \in X} \gamma(S, t) \neq \sum_{(S, t) \in X'} \gamma(S, t)$; w.l.o.g., assume that the left-hand side is larger. Then, if engine i were to increase the probability of the permutation σ by some $\epsilon > 0$ while reducing the probability of σ' by ϵ , her resulting change in utility would be

$$\begin{aligned} & \epsilon \sum_{(S, t) \in X} \gamma(S, t) f_i(\mathbf{q}(S, t, \mathcal{D})) - \epsilon \sum_{(S, t) \in X'} \gamma(S, t) f_i(\mathbf{q}(S, t, \mathcal{D})) \\ &= \epsilon \cdot f_i(1, \mathbf{0}) \cdot \left(\sum_{(S, t) \in X} \gamma(S, t) - \sum_{(S, t) \in X'} \gamma(S, t) \right) > 0. \end{aligned}$$

Engine i could thus increase her utility by shifting probability from σ' to σ , contradicting the equilibrium assumption. Thus, there must be another engine j with $q_j(S, t, \mathcal{D}_j) > 0$. Now, the non-indifference of the selection rules implies that $f_j(\mathbf{q}) < f_j(1, \mathbf{q}_{-j})$ for at least one such engine j (possibly $j = i$); for this particular engine j , this implies strictness of Inequality (5). The remainder of the proof now proceeds with this engine (formerly j) as engine i .

Finally, summing up the strict Inequality (5) over all types $(S, t) \in \text{Supp}(\Gamma)$ and changing orders of summation, we obtain that

$$u_i(\mathcal{D}) < \sum_{\sigma} \mathcal{D}_i(\sigma) \cdot \sum_{S \subseteq \Omega} \sum_{t \geq 0} g_i(\mathbb{1}[\sigma[t] \cap S \neq \emptyset]; S, t). \quad (6)$$

The right-hand side of Inequality (6) is exactly the utility of engine i when *mixing* over deterministic strategies instead of playing a non-deterministic strategy. Thus, \mathcal{D}_i cannot have been engine i 's best strategy at the correlated equilibrium under signal Z_i , a contradiction. \square

Proof of Theorem 4.1: Let $\mathbf{Q} = (\mathbf{q}_i)_{i=1}^K$ be a pure strategy equilibrium of the Singleton Search Engine Game. Assume for contradiction that \mathbf{Q} is not symmetric. Then, there are two engines (without loss of generality, engines 1 and 2) and two pages (without loss of generality, pages 1 and 2) such that $q_1(2) < q_2(2)$ and $q_1(1) > q_2(1)$. Let δ_1 be the constant in the definition of cross-concavity when applied to engines 1, 2 with $x = q_1(2)$ and $y = q_2(2)$. Similarly, let δ_2 be the constant when the definition is applied to engines 1, 2 with $x = q_2(1)$ and $y = q_1(1)$. Fix an arbitrary $\epsilon <$

$\min(\delta_1, \delta_2)$. By strict cross-concavity,

$$\begin{aligned} & f(q_1(2) + \epsilon, q_2(2), \mathbf{q}_{-\{1,2\}}) - f(q_1(2), q_2(2), \mathbf{q}_{-\{1,2\}}) \\ & > f(q_2(2), q_1(2), \mathbf{q}_{-\{1,2\}}) - f(q_2(2) - \epsilon, q_1(2), \mathbf{q}_{-\{1,2\}}), \\ & f(q_2(1) + \epsilon, q_1(1), \mathbf{q}_{-\{1,2\}}) - f(q_2(1), q_1(1), \mathbf{q}_{-\{1,2\}}) \\ & > f(q_1(1), q_2(1), \mathbf{q}_{-\{1,2\}}) - f(q_1(1) - \epsilon, q_2(1), \mathbf{q}_{-\{1,2\}}). \end{aligned} \quad (7)$$

Because \mathbf{Q} is an equilibrium, engine 1 cannot profit from moving ϵ probability from page 1 to page 2, nor can engine 2 profit from moving ϵ from page 2 to page 1. Thus,

$$\begin{aligned} & f(q_1(2) + \epsilon, q_2(2), \mathbf{q}_{-\{1,2\}}) + f(q_1(1) - \epsilon, q_2(1), \mathbf{q}_{-\{1,2\}}) \\ & \leq f(q_1(2), q_2(2), \mathbf{q}_{-\{1,2\}}) + f(q_1(1), q_2(1), \mathbf{q}_{-\{1,2\}}), \\ & f(q_2(1) + \epsilon, q_1(1), \mathbf{q}_{-\{1,2\}}) + f(q_2(2) - \epsilon, q_1(2), \mathbf{q}_{-\{1,2\}}) \\ & \leq f(q_2(1), q_1(1), \mathbf{q}_{-\{1,2\}}) + f(q_2(2), q_1(2), \mathbf{q}_{-\{1,2\}}). \end{aligned} \quad (8)$$

Adding the two Equations (7) and subtracting the two Equations (8), all terms cancel and we are left with the contradiction $0 > 0$. Hence, \mathbf{Q} must be symmetric. \square

Proof of Theorem 4.2: The proportional sharing rule satisfies symmetry and strict cross-concavity. Thus, by Theorem 4.1, all pure equilibria are symmetric. It remains to show that a symmetric equilibrium *exists* and is *unique*.

We begin by showing uniqueness and deriving a specific form for the equilibrium. Let \mathbf{Q} be a symmetric pure equilibrium; it is completely described by the (common) probability $q(n)$ with which the engines choose page n . We claim that $q(n) > 0$ for all pages n . Assume for contradiction that $q(n) = 0$ for some page n , and let n' be any page with $q(n') > 0$. Consider the deviation in which engine 1 adds a small amount ϵ of probability to page n , and removes it from page n' . Engine 1 receives all visits from users of type n , for an added payoff of $\gamma(n) > 0$ (since each page is desired with positive probability). The loss in payoff from page n' is less than $\gamma(n') \cdot \frac{\epsilon}{Kq(n')}$. For small enough ϵ , this is strictly less than the gain $\gamma(n)$. Thus, \mathbf{Q} was not an equilibrium.

Because $q(n) > 0$ for all n , the strategy profiles under consideration are bounded away in all coordinates from 0 (and also 1), and thus from the discontinuities (and non-differentiabilities) of the selection function. We can therefore safely take derivatives.

Engine i is choosing her strategy to maximize the value $\sum_n \gamma(n) f(\mathbf{q}(n))$ subject to $\sum_n q_i(n) = 1$; thus, the derivative of the Lagrangian $L(\mathbf{q}_i, \lambda) = \sum_n \gamma(n) f(\mathbf{q}(n)) + \lambda(1 - \sum_n q_i(n))$ with respect to all variables must be 0: that is, $\gamma(n) \frac{d}{dq_i(n)} f(q_i(n), \mathbf{q}_{-i}(n)) = \lambda$ for each page n . Using the definition of the proportional selection rule, this implies

$$\frac{\sum_{i' \neq i} q_{i'}(n)}{(\sum_{i'} q_{i'}(n))^2} = \frac{\lambda}{\gamma(n)} \quad (9)$$

for each n . As the equilibrium is symmetric, $q_{i'}(n) = q_i(n)$ for all i' , simplifying (9) to $q_i(n) = \frac{\gamma(n) \cdot (K-1)}{\lambda K^2}$. Since \mathbf{q}_i and Γ are distributions, we must have $\lambda K^2 / (K-1) = 1$, and hence $q_i(n) = \gamma(n)$ for all engines i is the only equilibrium.

It remains to prove that this distribution is in fact an equilibrium. Consider any best response \mathbf{q}_i by engine i to the symmetric equilibrium $q_{i'}(n) = \gamma(n)$ played by the other engines. We use variants of the Lagrangian condition (9), substituting that $q_{i'}(n) = \gamma(n)$. Assume $\mathbf{q}_i \neq \mathbf{q}_{i'}$; then there exist pages n, n' with $0 \leq q_i(n) < q_{i'}(n)$ and $q_{i'}(n') <$

¹⁰It would not have been strict if all indicator terms had been equal.

$q_i(n') \leq 1$. Since it is possible that $q_i(n) = 0$ or $q_i(n') = 1$, we only obtain weaker Lagrangian conditions of non-positive (resp., non-negative) derivative, as opposed to derivatives of 0. These weaker Lagrangian conditions for n, n' yield

$$\frac{(K-1)\gamma(n)}{(q_i(n) + (K-1)\gamma(n))^2} \leq \frac{\lambda}{\gamma(n)},$$

$$\frac{(K-1)\gamma(n')}{(q_i(n') + (K-1)\gamma(n'))^2} \geq \frac{\lambda}{\gamma(n')}.$$

We can rearrange for λ and use that $q_i(n) < \gamma(n)$ and $\gamma(n') < q_i(n')$ to obtain that

$$\frac{(K-1)\gamma(n)^2}{(K\gamma(n))^2} < \frac{(K-1)\gamma(n)^2}{(q_i(n) + (K-1)\gamma(n))^2}$$

$$\leq \lambda \leq \frac{(K-1)\gamma(n')^2}{(q_i(n') + (K-1)\gamma(n'))^2} < \frac{(K-1)\gamma(n')^2}{(K\gamma(n'))^2},$$

a contradiction. Thus, the only locally optimal response for player i is to play the same distribution as the other players, so the symmetric distribution is an equilibrium. \square

Proof of Proposition 4.3: Suppose that we have $N \geq K$ pages, and assume w.l.o.g. that $\gamma(1) > \gamma(2) > \dots > \gamma(N)$, and the probabilities are in general position. The social optimum occurs when engine i deterministically displays page i , leading to a social welfare of $\sum_{n=1}^K \gamma(n)$. By Theorem 4.2, the unique pure equilibrium has each engine choosing each page n with probability $\gamma(n)$. Each user will then be satisfied with probability $\sum_{n=1}^N \gamma(n)^2$, regardless of his engine choice. The PoS is thus $\frac{\sum_{n=1}^K \gamma(n)}{\sum_{n=1}^N \gamma(n)^2}$.

We now consider two special cases. First, when the first K pages have probability $\Theta(1/K)$ each, and the remaining pages together have probability $o(1/K)$, the PoS is $\Omega(K)$, as can be seen by upper-bounding $\sum_{n=1}^N \gamma(n)^2 \leq \sum_{n=1}^K \gamma(n)^2 + (\sum_{n>K} \gamma(n))^2$. Another distribution giving rise to this type of behavior is when all pages have probability $\Theta(1/N)$.

Second, consider the following distribution: the top two pages have probability $\Theta(1/\sqrt{N})$; the remaining pages have probability $\Theta(1/N)$. The PoS becomes $\Theta(\sqrt{N})$. Thus, when there are a few highly relevant pages and many rare pages, the symmetric equilibrium leads to inefficiency $\Omega(\sqrt{N})$. \square

Proof of Theorem 4.4: As in the proof of Theorem 4.2, if \mathbf{Q} is a symmetric equilibrium, say, $q_i(n) = q(n)$ for all i and n , then $q(n) > 0$ for all n ; otherwise, an engine could increase her utility by selecting page n with sufficiently small probability. We can therefore assume $q(n) > 0$ for all n .

Because each \mathbf{q}_i is a best response to the other engines' strategies, it maximizes

$$u_i(\mathbf{q}_i, \mathbf{q}_{-i}) = \sum_n \gamma(n) f(\mathbf{q}(n)) \cdot (\beta + (1-\beta)q_i(n))$$

$$= \sum_n \gamma(n) \frac{q_i(n)}{\sum_k q_k(n)} \cdot (\beta + (1-\beta)q_i(n)),$$

subject to $\sum_n q_i(n) = 1$. The Lagrangian for this constrained optimization, $L_i(\lambda_i, \mathbf{q}_i)$, is

$$\sum_n \gamma(n) \frac{q_i(n)}{q_i(n) + (K-1)q(n)} \cdot (\beta + (1-\beta)q_i(n))$$

$$+ \lambda_i (1 - \sum_n q_i(n)).$$

Because $q_i(n) > 0$ for all n , we can take derivatives; a necessary condition for \mathbf{q}_i to be a local optimum for L_i is then that $\frac{d}{dq_i(n)} L_i(\lambda_i, \mathbf{q}_i) = 0$ for each n . This yields

$$\gamma(n) \frac{(K-1)q(n) \cdot (2(1-\beta)q_i(n) + \beta) + (1-\beta)q_i(n)^2}{(q_i(n) + (K-1)q(n))^2} = \lambda_i.$$

Because \mathbf{Q} is symmetric, we also have that $q_i(n) = q(n)$ for all n , and can simplify this condition to

$$\gamma(n) \frac{(K-1)\beta + (2K-1)(1-\beta)q(n)}{K^2q(n)} = \lambda_i.$$

Writing $\lambda'_i = \frac{K^2\lambda_i}{(K-1)\beta}$, $z = \frac{(2K-1)(1-\beta)}{(K-1)\beta}$, and rearranging terms, we can solve for $q(n)$ to get

$$q(n) = \frac{\gamma(n)}{\lambda'_i - z\gamma(n)}. \quad (10)$$

Because $q(n) > 0$, this implies $\lambda'_i > z\gamma(n)$ for all n . The condition $\sum_n q(n) = 1$ implies

$$\sum_n \frac{\gamma(n)}{\lambda'_i - z\gamma(n)} = 1. \quad (11)$$

Now let $n^* \in \operatorname{argmax}_n \{\gamma(n)\}$ be a page maximizing the probability of being desired. For λ in the range $(z\gamma(n^*), \infty)$, we have $\frac{\gamma(n)}{\lambda - z\gamma(n)} > 0$ for all pages n , and the expression $\sum_n \frac{\gamma(n)}{\lambda - z\gamma(n)}$ is strictly decreasing, approaching 0 as $\lambda \rightarrow \infty$ and approaching ∞ as $\lambda \rightarrow z\gamma(n^*)$. There is therefore a unique solution (in λ'_i) to the Equation (11) in the range $(z\gamma(n^*), \infty)$. By choice of n^* , all terms of $\sum_n \frac{\gamma(n)}{\lambda'_i - z\gamma(n)}$ are positive at this solution, and each $\frac{\gamma(n)}{\lambda'_i - z\gamma(n)}$ lies in $(0, 1)$.

Thus, the $\frac{\gamma(n)}{\lambda'_i - z\gamma(n)}$ form a valid distribution over pages, and constitute the only local maximum of the Laplacian. Thus, the values $q(n) = \frac{\gamma(n)}{\lambda'_i - z\gamma(n)}$ constitute the only candidate symmetric equilibrium.¹¹

It remains to show that this symmetric profile, which attains a local maximum for the utility of each engine, is actually a global maximum (and thus a best response) for each player when $\beta > 1 - 1/K$. To see this, note that $\frac{d^2 u_i(q_i, \mathbf{q}_{-i})}{dq_i(n) dq_i(n')} = 0$ for all pairs of pages $n \neq n'$; furthermore, for each page n , $\frac{d^2 u_i(q_i, \mathbf{q}_{-i})}{d^2 q_i(n)}$ is equal to

$$-\frac{2\gamma(n)(\sum_{k \neq i} q_k(n)) \cdot (\beta - (1-\beta) \sum_{k \neq i} q_k(n))}{(\sum_k q_k(n))^3}. \quad (12)$$

Since $\sum_{k \neq i} q_k(n) \leq K-1$, our assumption that $\beta > 1 - \frac{1}{K}$ implies that $\beta - (1-\beta) \sum_{k \neq i} q_k(n) > 0$, and hence $\frac{d^2 u_i(q_i, \mathbf{q}_{-i})}{d^2 q_i(n)} < 0$ over the entire domain of $u_i(\cdot)$. This implies that the Hessian of $u_i(\cdot)$ is negative definite, and $u_i(\cdot)$ is strictly concave. Thus, the local maximum described above is also a global maximum. We conclude that the symmetric profile is an equilibrium. \square

¹¹We remark that for $\beta = 1$, we get that $z = 0$, so we recover the analysis from the proof of Theorem 4.2 as a special case.

B. REFERENCES

- [1] G. Aggarwal, J. Feldman, S. M. Muthukrishnan, and M. Pál. Sponsored search auctions with markovian users. In *Proc. 4th WINE*, pages 621–628, 2008.
- [2] E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In *Proc. 45th IEEE FOCS*, pages 295–304, 2004.
- [3] S. Athey and G. Ellison. Position auctions with consumer search. Technical Report 15253, National Bureau of Economic Research, Inc, August 2009.
- [4] C. d’Aspremont, J. J. Gabszewicz, and J.-F. Thisse. On hotelling’s ‘stability in competition’. *Econometrica*, 47(5):1145–50, September 1979.
- [5] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007.
- [6] H. Hotelling. Stability in competition. *The Economic Journal*, 39:41–57, 1929.
- [7] N. Immorlica, A. T. Kalai, B. Lucier, A. Moitra, A. Postlewaite, and M. Tennenholtz. Dueling algorithms. In *Proc. 42nd ACM STOC*, pages 215–224, 2011.
- [8] D. Kempe and M. Mahdian. A cascade model for externalities in sponsored search. In *Proc. 4th WINE*, pages 585–596, 2008.
- [9] R. Khoussainov. *Economics of Distributed Web Search: A Machine Learning Approach*. PhD thesis, Department of Computer Science, University College Dublin, 2004.
- [10] R. Khoussainov and N. Kushmerick. Distributed web search as a stochastic game. In *Proc. SIGIR 2003 Workshop on Distributed Information Retrieval*, pages 58–69, 2003.
- [11] J. M. Kleinberg and S. Oren. Mechanisms for (mis)allocating scientific credit. In *Proc. 42nd ACM STOC*, pages 529–538, 2011.
- [12] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proc. 17th STACS*, pages 404–413. Springer, 1999.
- [13] J. J. Liu and D. M. Chiu. Mathematical modeling of competition in sponsored search market. In *Proc. 2010 Workshop on the Economics of Networks, Systems, and Computation*, 2010.
- [14] A. Mas-Collel, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, 1995.
- [15] R. P. McAfee. Mechanism design by competing sellers. *Econometrica*, 61(6):1281–1312, November 1993.
- [16] T. Mukhopadhyay, U. Rajan, and R. Telang. The market structure for internet search engines. *Journal of Management Information Systems*, 21(2):137–160, 2004.
- [17] M. Peters and S. Severinov. Competition among sellers who offer auctions instead of prices. *Journal of Economic Theory*, 75(1):141–179, July 1997.
- [18] T. Roughgarden. Intrinsic robustness of the price of anarchy. In *Proc. 40th ACM STOC*, pages 513–522, 2009.
- [19] C. Shapiro and H. Varian. *Information Rules: A Strategic Guide to the Network Economy*. Harvard Business Press, 1999.
- [20] H. Varian. *Intermediate microeconomics: a modern approach*. W. W. Norton & Company, seventh edition, 2006.
- [21] A. Vetta. Nash equilibria in competitive societies with applications to facility location, traffic routing and auctions. In *Proc. 43rd IEEE FOCS*, pages 416–425, 2002.