

Search Design and Broad Matching

Kfir Eliaz and Ran Spiegler*

Tel-Aviv University, University of Michigan and University College London.

Abstract

We study decentralized mechanisms for allocating firms into search pools. The pools are created in response to noisy preference signals provided by consumers, who then browse the pools via costly random sequential search. Surplus-maximizing search pools are implementable in symmetric Nash equilibrium. Full extraction of the maximal surplus is implementable if and only if the distribution of consumer types satisfies a set of simple inequalities, which involve the relative fractions of consumers who like different products and the Bhattacharyya coefficient of similarity between their conditional signal distributions. The optimal mechanism can be simulated by a “keyword auction” with “broad matching”. (JEL D02, D47, D82).

Consider a consumer who has a certain need and looks for a product that will satisfy it. When confronted with a particular product, the consumer can identify whether he wants it. However, when he first embarks on the search process, he is unable to provide an exact description of his need, and can only submit an imprecise “query” that fits several product types. How should a benevolent planner react to the consumer’s query? If he gives the consumer a single item, he risks ending up with a poor match. Instead, the planner may provide the consumer with a *set of options* to browse. If search is time-consuming, the planner should design the set in order to balance two considerations: maximizing the probability that the consumer will find what he is looking for, and minimizing the amount of time it will take him to find it. In other words, the planner’s problem is to design the consumer’s “search environment”.

The consumer’s predicament is prevalent in other environments. When an employer approaches a Human Resource agency with an intention to hire a new worker, the most he can usually do is list a few characteristics that vaguely describe the kind of worker he is looking for. Likewise, when we look for a rental apartment through a real-estate agent, we typically describe only broad characteristics (location, size, amenities). A more modern example is online search, where users submit queries that often contain general and imprecise keywords. Imagine that you look for a specific piece of instrumental music on YouTube, but you forgot its name. You would recognize it instantaneously if you heard it, but the only information you can supply to YouTube is the piece’s genre. Finally, there are cases in which the consumer *passively* conveys information. For instance, in contemporary online platforms, the “cookies” on users’ computers describe their navigation history, which may be correlated with their current needs.

We attempt to capture situations of this kind with a simple model, in which every consumer is interested in only one type of product and can provide only one particular description or “query” concerning this product. Each product type is supplied by a continuum of firms; and when the consumer is given a subset of firms, he browses it using the most

basic search technology in the literature: random sequential search. In our model, a central planner aims to associate with each query a “search pool” (an infinite and typically heterogeneous collection of firms) so as to maximize expected social welfare, defined by the gross surplus created when consumers and suppliers transact minus the search costs that consumers incur. An efficient search pool is one in which it takes consumers little time on average to find what they want.

If consumers could provide a perfectly informative query, there would be no design issue: when a consumer says he is interested in product type x , the planner will respond with a homogenous search pool consisting of many suppliers of x . The design problem arises when the consumer’s query is a noisy signal of his need. The planner could still reach his objective if he could directly identify the type of product offered by each supplier: the optimal composition of product types in the search pool associated with each query is determined by a simple first-order condition that reflects the preference distribution among consumers who submitted the query.

The problem becomes economically interesting when the planner wishes to *decentralize* the allocation of suppliers into search pools (because it is too costly for him to directly verify product types, or because he wants to automatize the allocation procedure). In this case the planner needs to make sure that suppliers’ assignment into search pools is aligned with their incentives. This is an object-allocation problem, where the allocated objects are unusual: entitlements to enter consumers’ search pools. A supplier’s ultimate evaluation of this allocation is a function of the joint distribution of consumers’ needs and queries (as well as the assignment of other suppliers). The incentive problem arises because consumers’ queries are noisy signals of their needs, which encourages suppliers to present themselves as relevant to many queries. In the absence of appropriate design, consumers’ search pools will be inefficiently composed, resulting in excessively prolonged search.

Our first result is that as long as the planner has no budget constraints, there is a direct (anonymous) mechanism that Nash implements the efficient search pools for *any* underlying

joint distribution over consumers’ needs and queries. Next, we ask whether the planner can design a direct mechanism that maximizes social surplus *and* fully extracts it. If the planner were a perfect monopolist trying to maximize profits, this would be his first-best outcome. This objective turns out to be implementable if and only if every pair of product types satisfies a simple inequality, which incorporates two quantities: the product types’ *relative popularity*, and a conventional measure (known as the Bhattacharyya Coefficient) of the *similarity between the query distributions* that characterize consumers who look for these two product types. As consumer preferences become more uniformly distributed and queries become more Blackwell-informative of consumers’ preferences, the perfect monopolist’s objective is more easily implementable. When it is implemented, the “price per impression” that firms of any given type pay obeys a simple formula. These prices decrease with the Blackwell informativeness of consumers’ queries.

Can the optimal mechanism be mimicked by an indirect mechanism in which suppliers engage in competitive bidding for the right to be included in search pools? In other words, *can some form of a competitive market implement efficient search environments?* We propose an auction format in which each firm submits a “bid-per-impression” for one query of its choice, and where the highest bidder for one query can enter with some probability into search pools associated with other queries that it did not bid for. We refer to this last feature as “broad matching”, hijacking search-engine jargon. Thus, when a consumer submits a query, he receives a potentially mixed search pool that includes firms that bid for other queries. We show that when the perfect monopolist’s objective is implementable, our broad-match auction has an essentially unique symmetric pure-strategy Nash equilibrium, in which social surplus is maximized and fully extracted. Furthermore, the optimal auction has a simple specification of the “broad match function” that determines the probability with which the highest bidders for a query enter the search pool associated with any given query.

Our notion of broad matching captures an intuitive function of real-life institutions that match agents in two-sided markets. This function can be described as “vocabulary expan-

sion”. When a prospective buyer looks up “road bikes” in a classified directory, he may see items listed by sellers under “racing bikes”, “hybrid bikes” or “fixed gear bikes”. Although the buyer and seller used different terms, the directory bridges the gap between their vocabularies in the interest of a potentially good match. Likewise, “organic” online search engines respond to keyword-based queries by taking into account typos and semantically related terms. The broad-match function plays a similar role in our indirect mechanism.

The main contributions of this paper are thus threefold: formulating the problem of decentralized implementation of efficient search environments, elucidating the forces that shape the relevant incentive constraints, and the novel broad-matching auction format. Although this project has been inspired by modern online search engines, the model obviously abstracts from some of their key features and therefore cannot be viewed as a faithful model of this real-life institution. We do hope, however, that our abstraction will help researchers conceptualize the problem of “search design” in various economic settings, which include search engines, but - as we demonstrate in the concluding section - are not restricted to them.

I. The environment

Let X be a finite set of *product types*, $|X| \geq 2$. Let W be a finite set of *signals*, where $|W| \geq |X|$. For every $x \in X$, there is a measure one of *firms* that offer only this product type (as many units as required, at zero cost). We sometimes refer to them as x firms. Each firm is informed of its own type only. A firm gets a fixed payoff of 1 from any unit it sells (we abstract from product prices).¹

On the other side of the market, there is a measure one of *consumers*. A consumer type is defined by the pair (x, w) , where x is the (only) type of product he is interested in, and w is the signal he can convey about what he is looking for. In line with some of the potential applications we have in mind, we often refer to w as a “query”, a “keyword” or “the consumer’s vocabulary”. Let $\mu \in \Delta(X \times W)$ be the distribution of consumer types

in the population. The marginals of μ on X and W have full support. As usual, denote $\mu(x) = \sum_w \mu(x, w)$ and $\mu(\cdot | x) = (\mu(w | x))_{w \in W}$. The latter is referred to as the conditional signal/query distribution that characterizes the preference type x .

When a consumer of type (x, w) consumes a product of type $y \neq x$, he gets a sure payoff of 0. When he consumes a particular product of type x , he gets a payoff of $\pi > 0$ with independent probability θ and a payoff of zero with probability $1 - \theta$. The parameter θ captures idiosyncratic heterogeneity among consumers and firms of a given type. Products are “inspection goods”: when a consumer encounters a particular product, he immediately recognizes the payoff it generates. However, all he can communicate ex-ante about what he is looking for is encapsulated in the signal w . Note that w does *not* represent the consumer’s information about his own preferences, but the information he can provide to *others*. Thus, when the consumer inspects a particular product, this does not cause him to revise his beliefs about other products.

Consumers decide *ex-ante*, before their type is realized, whether to engage in search. If a consumer opts out, he gets a payoff of zero. If a consumer decides to search, his type (x, w) is subsequently realized and he automatically submits the signal/query w . He does not get to revise his search decision in this interim stage. In response to his query, the consumer is given access to an infinite collection of products. We refer to this collection as the *search pool* associated with w , and define it formally as a probability distribution $(\lambda(x | w))_{x \in X}$, where $\lambda(x | w)$ is the fraction of products of type x in the pool. Denote $\lambda = (\lambda(\cdot | w))_{w \in W}$. The consumer repeatedly draws independent random samples from this pool (with replacement), where each draw carries a fixed cost $c \in (0, \theta)$. As soon as the consumer finds a product that satisfies his need, he transacts with its seller and terminates his search. If $\lambda(x | w) > 0$, a consumer of type (x, w) will find what he wants in finite time with probability one (because the pool contains infinitely many products).

The consumer’s ex-ante participation decision maximizes his expected utility according to "rational expectations". In fact, all we need to assume is that the consumer’s decision

is based on a correct estimate of: (1) the probability he will find a product that satisfies his (yet to be realized) need; and (2) the expected duration of his search. These are broad features of the overall quality of a search platform, and consumers can plausibly learn them in the long run.

An example: Mozart vs. Stravinsky

The following specification illustrates the primitives of our model and serves as a running example. Our consumer population consists of people who had a previous chance encounter with some piece of classical music - while listening to the radio, watching a film or attending a public ceremony - and are now trying to retrieve it (different consumers are looking for different pieces). Suppose that all relevant pieces have been composed by either Mozart or Stravinsky. Accordingly, let $X = \{moz, str\}$. The set of signals consists of three keywords, “Mozart”, “Stravinsky” and “Classical Music”, denoted MOZ , STR and CL respectively. A consumer type (moz, MOZ) ((str, STR)) is interpreted as someone who recognizes that the piece he once heard was by Mozart (Stravinsky). In contrast, the types (moz, CL) and (str, CL) are unable to describe the composer of the piece they are looking for, and all they can say ex-ante is that they are looking for some piece of classical music. Each consumer can recognize whether a particular music file he listens to is a rendition of the piece he is looking for. However, the inspection never reveals the piece’s composer, and consequently the consumer is never able to revise his query. Assume $\mu(moz, STR) = \mu(str, MOZ) = 0$ - that is, when a consumer can name the composer, he never makes an error.

Optimal search pools

Consider a central planner who aims to allocate firms into search pools in order to maximize total surplus. Since consumers make their participation decisions before their types are realized, we can assume that they either all participate or all opt out. Conditional on participation, social surplus is given by the following function of the collection of search pools λ :

$$U(\lambda) \equiv \sum_{(x,w)|\lambda(x|w)>0} \mu(x, w) \cdot (1 + \pi) - \sum_{(x,w)} \mu(x, w) \cdot \frac{c}{\theta \cdot \lambda(x | w)}$$

It is clear from this expression that the planner will set $\lambda(x | w) > 0$ whenever $\mu(x, w) > 0$, to prevent a positive mass of consumers from incurring infinite search costs. Thus, conditional on participating, all consumers eventually find what they want, such that firms earn a total payoff of 1, consumers earn a total gross payoff of π , and the expected search cost of a consumer of type (x, w) is $c/\theta\lambda(x | w)$.

Define $\lambda^* = \arg \max_{\lambda} U(\lambda)$. This means that λ^* minimizes consumers' expected search duration. First-order conditions imply

$$(1) \quad \frac{\lambda^*(x | w)}{\lambda^*(y | w)} = \sqrt{\frac{\mu(x, w)}{\mu(y, w)}}$$

whenever $\mu(x, w)\mu(y, w) > 0$. Since $\sum_{x \in X} \lambda^*(x | w) = 1$, we obtain

$$(2) \quad \lambda^*(x | w) = \frac{\sqrt{\mu(x, w)}}{\sum_{y \in X} \sqrt{\mu(y, w)}}$$

In the ‘‘Mozart vs. Stravinsky’’ example, $\lambda^*(str, STR) = \lambda^*(moz, MOZ) = 1$, because the signals *MOZ* and *STR* are perfectly informative of the consumer's preferences. The fraction of *moz* products in the search pool associated with the query *CL* is

$$(3) \quad \lambda^*(moz, CL) = \frac{\sqrt{\mu(moz, CL)}}{\sqrt{\mu(moz, CL)} + \sqrt{\mu(str, CL)}}$$

From now on, we will assume that c is sufficiently small, such that $U(\lambda^*) > 0$ - i.e., it is socially optimal for consumers to participate.

Discussion

Before turning to the analysis, we comment on the way we model consumer search. Since this paper is the first to formulate the ‘‘search design’’ problem, we were guided by simplicity, and opted for the simplest (and arguably most common) search protocol in the literature: random sequential sampling from an infinitely large pool of items. Important applications of this protocol to consumer-market settings include Wolinsky (1986) and Armstrong et al.

(2009). Thanks to this simplification, we can afford not to sacrifice generality in modeling the novel friction at the heart of our paper, namely consumers' limited ability to describe their needs. Modifications of the search protocol that assume large finite pools or introduce a small exogenous stopping probability (capturing "search fatigue") would lead to very similar, but less elegant characterizations.

The *random-sampling* assumption means that the planner is unable to control the order by which consumers inspect products in their search pool. This is a reasonable assumption for search environments in which inspection is done "offline". For instance, when a consumer obtains a list of specialists from a "Yellow Pages" directory, he can verify whether a specialist is a good match only by physically contacting him. The order in which he will examine the specialists will depend on their availability, which lies beyond the control of the directory's planner. Even in the case of online search - where search results are ordered - the platform's ability to control the order of inspection is imperfect. Athey and Ellison (2011) list a number of reasons why web users may disobey the order in which links appear on their computer screen: advertisers attract their attention away from prominent links; some links are slow or broken; and they may distrust the motives behind the search engine's suggested order. From this point of view, our random-search assumption fits a *worst-case* analysis for the design of search environments. We further discuss ordered search in Section 5.

Taking the random-sampling assumption as given, the social value of having *infinitely many* firms in search pools arises from our twin assumptions that $\theta < 1$ and that the sampling process is with replacement. If we assumed that $\theta = 1$, there would be no need for search pools to include more than one firm of each type. Likewise, if the random sampling procedure were *without* replacement (an admittedly more realistic assumption), one could at least construct examples in which a one-of-a-type pool dominates an infinite pool. However, an infinite pool is socially optimal when sampling is with replacement.

Finally, consider our restriction that consumers make a single, ex-ante participation decision. As long as π is sufficiently large, it is easy to extend our analysis to the case in

which the consumer can revise his search decision after his type is realized (in particular, he can choose to stop searching at any point); the results are essentially the same (the difference is that efficient search pools may now exclude consumer and firm types). A deeper challenge is that in reality, query-based search typically proceeds via successive alterations of the initial query. The reason is that after encountering several items in the search pool, consumers get a better sense of how they ought to describe their need, and consequently revise their query. Constructing richer models of query-based search that incorporate this idea is an important project for future research, even outside the context of “search design”.

II. Mechanism Design

We are interested in situations where only the consumer can verify whether a particular product satisfies his need. Moreover, the only means of verification is personal inspection via sequential search. In these situations, the planner cannot monitor the composition of search pools. In order to design search pools with a particular composition, firms need to be incentivized to enter the appropriate search pools.

A mechanism requires each firm to choose an element from some message space (the firm can also opt out, in which case it earns 0). For every profile of messages, the mechanism allocates firms probabilistically to the search pools associated with each $w \in W$, and specifies a (possibly negative) transfer that the firm pays to the planner. In addition, the mechanism specifies a (possibly negative) transfer to consumers conditional on participation. The mechanism induces a simultaneous-move game, in which consumers choose (before their types are realized) whether to participate, and each firm submits its message or opts out. Following this game, consumers’ types are realized and they automatically submit their queries (if they chose to participate), and the search pool associated with each query is induced by the mechanism and firms’ strategies. The solution concept we use is *symmetric Nash equilibrium* - i.e., all consumers play the same strategy, and all firms of a given type play the same strategy.

In this section, we focus on *direct* mechanisms - i.e. the firms' message space is X . We also restrict attention to *anonymous* mechanisms: the outcome for each firm is purely a function of its own report and the overall distribution of firms' reports.

What can the planner monitor?

Our definition of direct mechanisms specifies lump-sum transfers that are independent of events that unfold inside each search pool. Taken literally, this fits situations in which the planner can only monitor whether firms *access* search pools. This is the case of a "Yellow Pages" directory, for example. However, our specification is also consistent with the case in which the planner can monitor "*impressions*" - i.e., he can condition the firm's transfer on the *number of draws* it receives in each search pool (online search engines fall into this category). To see why, note that in our model firms are risk-neutral, hence all they care about is the total number of transactions minus the total payment induced by any message they submit. Therefore, our reduction to lump-sum transfers is a mere simplification that carries no loss of generality.

Things are different when the planner can monitor *transactions*. This assumption is sensible in environments like real-estate intermediation, where transactions are verifiable and subjected to exclusive-dealership arrangements. In our model, it would completely trivialize the planner's problem. By assumption, all firms earn a payoff of 1 conditional on a transaction, regardless of the transacting parties' types. Therefore, the planner could simply set a uniform price-per-transaction of 1, and thus implement any outcome he wishes. It follows that the interest in our problem arises only when the planner is unable to monitor transactions.

A. Implementing Efficient Search Pools

In this sub-section we assume that the planner's sole objective is to maximize social surplus. He has an unlimited budget and can therefore satisfy the participation constraints of consumers and firms. By the Revelation Principle, we assume w.l.o.g that conditional on

participating, each firm truthfully reports its type in equilibrium. In addition, we restrict attention to equilibria in which all firms choose to participate with probability one (we show in the Appendix why this restriction is w.l.o.g). We are now able to reduce any anonymous direct mechanism to the pair $\langle q, T \rangle = \langle (q(x, w))_{x \in X, w \in W}, (T_x)_{x \in X} \rangle$, such that an individual firm that reports \hat{x} enters the search pool associated with w with probability $q(\hat{x}, w)$, and pays the transfer $T(\hat{x})$ to the planner.² Because consumers have no private information at the time they make their decision, we omit their transfer from the description and take it for granted that it ensures their participation.

The planner's problem is to find an anonymous direct mechanism that sustains the optimal collection of search pools λ^* in symmetric Nash equilibrium. By the assumption that all firms participate and report truthfully, the allocation rule q must satisfy

$$(4) \quad \lambda^*(x | w) = \frac{q(x, w)}{\sum_y q(y, w)}$$

for every (x, w) in the support of μ . It follows that the planner's problem is to find a pair $\langle q, T \rangle$ that induces λ^* according to (4), subject to the firms' *incentive compatibility* constraint: for every $x, y \in X$,

$$(5) \quad \sum_{w \in W} q(x, w) \cdot \frac{\mu(x, w)}{q(x, w)} - T_x \geq \sum_{w \in W} q(y, w) \cdot \frac{\mu(x, w)}{q(x, w)} - T_y$$

To understand this inequality, recall that when a firm of type x submits the report \hat{x} , it enters the search pool associated with any query w with probability $q(\hat{x}, w)$. Let us calculate the number of transactions the firm expects in the pool. The total measure of consumers at the pool who are interested in the firm's type of product is $\mu(x, w)$. All of these consumers eventually find a desirable product; hence, they are equally shared by the x firms in the pool. The total measure of these firms is $q(x, w)$, by the assumption that all firms participate and report truthfully. The number of transactions that the firm expects to get in the pool associated with w is thus $\mu(x, w)/q(x, w)$. The firm's total number of

transactions is then calculated by summing over all consumer queries. To get the firm's payoff, we subtract the transfer $T_{\hat{x}}$ induced by the firm's report. Note that the firm's net equilibrium payoff is the L.H.S of (5), which is equal to $\mu(x) - T_x$.

Plugging (1) and (4) into (5) and rearranging, we obtain the following lemma.

Lemma 1. *An anonymous direct mechanism defined by $\langle q, T \rangle$ implements λ^* in a symmetric Nash equilibrium if and only if*

$$(6) \quad \mu(x) - \sum_{w \in W} \sqrt{\mu(x, w)\mu(y, w)} \geq T_x - T_y$$

for every $x, y \in X$.

We are now ready for our first main result.

Proposition 2. *There is an anonymous direct mechanism that implements λ^* in symmetric Nash equilibrium.*

Proof. We prove a stronger result: the first-best is implementable by *independent* mechanisms for each w , as if W consisted of w alone. This result is stronger because we continue to assume the same equilibrium behavior by firms (participation and truthful revelation) and the same implemented outcome, but there are more constraints to satisfy (separate participation and IC constraints for each signal, compared with single participation and IC constraints in the grand mechanism).

Consider some $w \in W$ and some pair of product types $x, y \in X$. When the relevant set of signals is a singleton, $\{w\}$, the constraint that prevents type x from pretending to be y , denoted $IC(x, y)$, is given by the following inequality (derived from (6)):

$$(7) \quad \mu(x, w) - \sqrt{\mu(x, w)\mu(y, w)} \geq T_x - T_y$$

Let $\phi(x, y)$ denote the L.H.S. of (7), and rewrite $IC(x, y)$ as $\phi(x, y) \geq T_x - T_y$. For any cycle

of products $(x_1, x_2, \dots, x_m, x_1)$,

$$\begin{aligned} \phi(x_1, x_2) + \dots + \phi(x_m, x_1) &= \sum_{i=1}^m \left(\mu(x_i, w) - \sqrt{\mu(x_i, w)\mu(x_{(i+1) \bmod m}, w)} \right) \\ &\geq \sum_{i=1}^m \left(\mu(x_i, w) - \frac{\mu(x_i, w) + \mu(x_{(i+1) \bmod m}, w)}{2} \right) \\ &= 0 \end{aligned}$$

Then, by Rochet (1987), there is a collection of transfers $(T_x)_{x \in X}$ that satisfies $IC(x, y)$ for all $x, y \in X$.

To illustrate the argument underlying the proof, let $X = \{x, y\}$, $W = \{w\}$. Suppose the function q is designed to induce an efficient search pool under truthful reports. If type x reports truthfully, he gets $\mu(x)$ transactions. If he misreports, then by Lemma 1, the number of transactions he will get is the geometric average of $\mu(x)$ and $\mu(y)$. The mutual incentive compatibility constraints are then reduced to the requirement that the arithmetic average of $\mu(x)$ and $\mu(y)$ is weakly greater than their geometric average, which is always true. In the Appendix, we explicitly derive transfers that implement λ^* , using Vohra's (2011) flow-network technique.

B. Full Surplus Extraction and Bhattacharyya Similarity

In this sub-section we assume that the planner does not settle for maximizing social surplus, but also wants to fully extract it. If the planner were a monopolistic, profit-maximizing search platform, that would be its ideal outcome. Accordingly, we refer to the planner's goal in this sub-section as the "*perfect monopolist's objective*". We focus on the firms' behavior and take it for granted that the planner sets consumers' transfer such that their participation constraint holds bindingly.

Before characterizing the implementability of this more stringent objective, we need to introduce a new concept. When a firm contemplates misreporting its type, it realizes that it

may be allocated to search pools associated with different queries. However, since consumers' queries are noisy signals of their needs, the firm may still get a fair number of transactions despite its misallocation. Thus, the firm's potential gain from misreporting depends on the distinctiveness of the signal distributions that characterize consumer preference groups. It is therefore useful to measure the similarity between the conditional query distributions $\mu(\cdot | x)$ and $\mu(\cdot | y)$, for any given x and y . Of course, there are multiple ways to measure similarity between probability distributions. However, one particular measure turns out to be relevant for our exercise. For any pair of products $x, y \in X$, define

$$S(x, y) \equiv \left(\sum_{w \in W} \sqrt{\mu(w | x)\mu(w | y)} \right)^2$$

In the statistics literature, $\sqrt{S(x, y)}$ is known as the *Bhattacharyya Coefficient* that characterizes the distributions $\mu(\cdot | x)$ and $\mu(\cdot | y)$.³ From a geometric point of view, this is an appropriate similarity measure, because $\sqrt{S(x, y)}$ is the direction cosine between two unit vectors $(\sqrt{\mu(w | x)})_{w \in W}$ and $(\sqrt{\mu(w | y)})_{w \in W}$. The value of $S(x, y)$ increases as the angle between these two vectors shrinks; $S(x, y) = 1$ (0) if the two vectors are identical (orthogonal). More importantly, the following observation establishes that $S(x, y)$ is an appropriate similarity measure given our interpretation of $\mu(\cdot | x)$ and $\mu(\cdot | y)$ as conditional signal functions.

Remark 3. *Regard the stochastic matrix $(\mu(\cdot | x))_{x \in X}$ as an information system in Blackwell's sense. When $(\mu(\cdot | x))_{x \in X}$ is subjected to Blackwell garbling, $S(x, y)$ weakly increases for all x, y .*

Thus, Bhattacharyya Coefficients decrease with the Blackwell-informativeness of consumers' queries. This is consistent with the intuition that the distance between conditional query distributions captures consumers' ability to describe their needs. In our "Mozart vs. Stravinsky" example, $S(moz, str) = \mu(CL | moz)\mu(CL | str)$.

Our next result establishes a necessary and sufficient condition for the implementability

of the perfect monopolist's objective.

Proposition 4. *There exists an anonymous direct mechanism that implements the perfect monopolist's objective if and only if*

$$(8) \quad \frac{\mu(x)}{\mu(y)} S(x, y) \leq 1$$

for every $x, y \in X$.

Proof. As pointed out in the discussion below (5), the net payoff of a firm of type x when all firms report truthfully is $\mu(x) - T_x$. Consider any $x, y \in X$. Full surplus extraction means that $\mu(x) - T_x = \mu(y) - T_y = 0$. Thus, $IC(x, y)$, as given by (6), is reduced to $\mu(y) \geq \sqrt{\mu(x)\mu(y)S(x, y)}$, which is equivalent to (8).

Thus, the forces that obstruct the perfect monopolist's objective are large popularity gaps between products and uninformative signals. Consider two consumer preference types, x and y . If x is significantly more common than y in the consumer population, and the conditional signal distributions that characterize x and y are relatively similar, then $IC(x, y)$ will fail to hold. In the "Mozart vs. Stravinsky" example, $IC(moz, str)$ is violated when $\mu(moz)/\mu(str)$ is far from 1 and $\mu(CL)$ is large. At the other extreme, the perfect monopolist's objective is implementable when the marginal of μ over X is uniform (because $S(x, y) \leq 1$ for every x, y), or when consumers' signals are fully informative (because in this case, $S(x, y) = 0$ for all $x \neq y$).

To get an intuition for condition (8), let $X = \{x, y\}$ and $\mu(x) > \mu(y)$, and let firms play an equilibrium that sustains the first-best. Now suppose that a firm of type x deviates by pretending to be y . On one hand, the firm's transfer to the planner drops from $\mu(x)$ to $\mu(y)$. On the other hand, the firm's probability of entering the search pool associated with any w drops less in relative terms, because by (1) and (4), the ratio $q(y, w)/q(x, w)$ is equal to the *square root* of the ratio $\mu(y, w)/\mu(x, w)$. Thus, there is a sense in which, relative to the transfer rule, the allocation rule is biased in favor of the minority preference group, and this

boosts the x firm's incentive to deviate. If consumers' signals are sufficiently noisy, there will be many consumers who want x in the search pools to which the firm is misallocated, and $IC(x, y)$ will be violated.

We saw that when the planner only cares about efficiency, signal-by-signal implementation of his objective is possible. In contrast, for generic μ , signal-by-signal implementation of the perfect monopolist's objective is impossible. The reason is that in a single-query environment, $S(x, y) = 1$ for all x, y . Hence, Inequality (8) is violated by product types x, y for which $\mu(x) > \mu(y)$.

Prices per impression

The price-per-impression that a firm pays is the total transfer the firm pays to the planner, divided by the total number of times it is drawn by consumers in search pools. In any equilibrium that implements the perfect monopolist's objective, the price-per-impression of a firm is equal to its average "conversion rate" - i.e., the ratio of the total number of transactions the firm makes to the total number of draws it gets. The next result characterizes this quantity.

Proposition 5. *A direct anonymous mechanism that implements the perfect monopolist's objective induces the following price-per-impression for every product type x :*

$$(9) \quad p^*(x) = \frac{\theta \sqrt{\mu(x)}}{\sum_y \sqrt{\mu(y) S(x, y)}}$$

Bhattacharyya Coefficients play an important role in this price function: $p^*(x)$ increases as the conditional query distribution that characterizes consumers who look for x becomes more distinctive. The price-per-impression of all x is weakly increasing in the Blackwell informativeness of consumers' queries.

To illustrate the comparative statics, suppose that $X = W$ and that the marginal of μ over X is uniform. Consider two extreme cases. First, let $\mu(x | x) = 1$ for all x - i.e., consumers can perfectly describe their needs. Then, $p^*(x) = \theta$ for every x . Second, suppose

μ is uniform over $X \times X$. In this case, consumers' signals are entirely uninformative, and we have $p^*(x) = \theta/|X|$ for every x .

Finally, in the ‘‘Mozart vs. Stravinsky’’ example,

$$\begin{aligned} p^*(str) &= \frac{\theta}{1 + \sqrt{\mu(moz, CL)/\mu(str, CL)}} \\ p^*(moz) &= \frac{\theta}{1 + \sqrt{\mu(str, CL)/\mu(moz, CL)}} \end{aligned}$$

Note that $p^*(moz)$ ($p^*(str)$) is an increasing (decreasing) function of the fraction of *moz* fans among consumers who submit the query CL .

C. Uniform-Price Mechanisms

Suppose that the planner is not interested in extracting surplus. Instead, he wishes to implement λ^* using a mechanism in which transfers are independent of the firms' reports - i.e., $T_x = T_y$ for all $x, y \in X^*$. In particular, the mechanism may be required to involve *no transfers at all*. We refer to anonymous direct mechanisms that satisfy this additional requirement as uniform-price mechanisms.

Proposition 6. *There exists a uniform-price mechanism that implements λ^* in symmetric Nash equilibrium, if and only if (8) holds for every $x, y \in X$.*

Proof. Let $x, y \in X$. Impose the uniform-price requirement $T_x = T_y$. Then, $IC(y, x)$, as given by (6), is reduced to $\mu(y) \geq \sqrt{\mu(x)\mu(y)S(x, y)}$, which is equivalent to (8). (Note that the result follows from examining $IC(y, x)$, whereas in the proof of Proposition (4) it followed from examining $IC(x, y)$.)

Thus, the condition for implementability of efficient search pools by a uniform-price mechanism is the same as the condition for implementability of the perfect monopolist's objective. To appreciate this coincidence, consider an abstract single-agent mechanism-design setting with two agent types, 1 and 2, and two possible outcomes, O_1 and O_2 . Suppose

that the planner’s objective is to assign the outcome O_k to type k . The gross payoff of type i from outcome O_j is denoted d_{ij} . Let t_k denote the transfer the agent pays when he reports his type to be k . The IC constraints that ensure truthful reporting are $d_{11} - t_1 \geq d_{12} - t_2$ and $d_{22} - t_2 \geq d_{21} - t_1$. Full surplus extraction implies $t_k = d_{kk}$, and the IC constraints are reduced to $d_{22} \geq d_{12}$ and $d_{11} \geq d_{21}$. Alternatively, uniform prices mean $t_1 = t_2$, and the IC constraints are reduced to $d_{11} \geq d_{12}$ and $d_{22} \geq d_{21}$. Obviously, these systems of inequalities need not coincide. And indeed, in general settings, the conditions for implementing an outcome with full surplus extraction and with a uniform-price mechanism are different. However, if d_{12} and d_{21} happen to be the same, the two systems do coincide. This is precisely what happens in our setting: we saw that under λ^* , the number of transactions that a firm of any type x expects when it pretends to be any other type y is $\sum_{w \in W} \sqrt{\mu(x, w)\mu(y, w)}$.

III. Competitive Bidding and Broad Matching

Is there a plausible *indirect* mechanism that implements the perfect monopolist’s objective when condition (8) holds for every $x, y \in X$? We are particularly interested in mechanisms based on *competitive bidding*, because they are natural benchmarks for object-allocation mechanisms and commonly used by online search engines. Moreover, they shed light on a key question: can “competitive forces” generate efficient search environments?

We propose a mechanism referred to as a “*broad-match auction*”. Each firm simultaneously chooses exactly *one* signal in W to bid for. Bids are “per impression”, according to a rule we explain below. We introduce a *broad-match function* $b : W \times W \rightarrow [0, 1]$ that assigns to each signal $v \in W$ a measure $b(w | v)$ of “ (w, v) tickets” for every $w \in W$. Each of these (w, v) tickets is uniformly assigned to the firms that submitted the highest bid for v . When a consumer submits the query w , he gets access to a search pool consisting of a measure $\sum_{v \in W} b(w | v)$ of (w, \cdot) tickets, originating from all signals v via the broad-match function. The consumer then sequentially samples random *tickets* from this pool. Each time a consumer draws a (w, v) ticket, he incurs a search cost c and inspects the firm holding the

ticket, and the firm pays to the planner the (winning) bid-per-impression it submitted for v . As to the consumer's transfer, assume it makes him indifferent to participation under the efficient allocation of firms to search pools.

From consumers' point of view, the broad-match function plays the role of "vocabulary expansion". When a consumer submits the query w , he receives a search pool consisting of "tickets" that originate from various signals v . Therefore, the broad-match function mimics an environment in which the consumer has a richer vocabulary. Real-life intermediaries regularly fulfill this role, in order to generate matches that otherwise would not occur due to the different language that parties on the two sides of the market use to describe themselves or what they are looking for. Our innovation will be to show that when this role is *explicitly integrated into the design of the search platform*, it can augment a conventional competitive-bidding mechanism and go beyond its limitations.

The search process under the broad-match auction is slightly different from the description in Section 2, because here consumers repeatedly sample *tickets* rather than suppliers. The *conversion rate* that a (w, v) ticket generates for the individual firm that holds it is the probability that a consumer who draws the ticket will transact with that firm. The average conversion rate that the signal v generates for an individual firm that submitted the highest bid for v is the probability that a consumer who draws *any* of the (\cdot, v) tickets the firm holds will transact with it.

The broad-match auction admits the possibility that all tickets in a search pool will be held by a single firm. This possibility will not be realized in the equilibria we will analyze, but it can occur off the equilibrium path - e.g. when a single firm outbids all others for some signal. The firm's payoff in this case is ill-defined because it receives infinitely many draws. All we need to assume in this case is that if the conversion rate that the signal generates for that firm is above (below) the firm's bid, the firm's payoff is arbitrarily high (low).

As a benchmark for our analysis, let us consider two extreme specifications of b .

Exact matching

Suppose that $b(w | w) = 1$ and $b(w | v) = 0$ for all $w \neq v$. In this case, we have an “exact-match auction”: consumers who submit the query w are brought into contact only with firms that submitted a winning bid for w . This reduces our mechanism to a signal-by-signal first-price auction. If consumers’ signals were perfectly informative, it could implement the perfect monopolist’s objective in symmetric pure-strategy Nash equilibrium. The difficulty with exact matching arises when signals are noisy, such that optimal search pools are required to be diverse. Consider the “Mozart vs. Stravinsky” example. If we restrict attention to symmetric pure-strategy profiles, exact matching only allows $\lambda(\cdot | CL)$ to take the values $0, \frac{1}{2}, 1$, and therefore it is mechanically unable to generate $\lambda^*(\cdot | CL)$ for generic μ .

However, the difficulty with exact matching extends to mixed-strategy equilibria. Let $m(x)$ denote the measure of x firms in the CL search pool that is induced by some symmetric (and possibly mixed) Nash equilibrium. Then,

$$(10) \quad \frac{m(moz)}{m(str)} = \frac{\lambda(moz | CL)}{\lambda(str | CL)}$$

In order for λ to coincide with λ^* , it must be the case that $m(moz) > 0$ and $m(str) > 0$ - i.e., both moz and str firms must submit the winning bid for CL . In order for the firms’ surplus to be fully extracted, the winning bid must be equal to the conversion rate they expect at CL . Therefore, they must expect the same number of transactions in the pool. This number for an individual x firm is $\mu(x, CL)/m(x)$, because all consumers of type (x, CL) search and eventually find a product they like, and all x firms in the pool equally share this clientele. Plugging (10), we obtain

$$\frac{\lambda(moz | CL)}{\lambda(str | CL)} = \frac{\mu(moz, CL)}{\mu(str, CL)}$$

which contradicts the efficiency condition (3).

Fully broad matching

Suppose that $b(w | v) = 1$ for all w, v . This reduces our mechanism to exact matching,

defined for a fictitious environment that consists of one signal w^* , such that the fraction of (x, w^*) consumers is $\mu(x)$. Therefore, fully broad matching is weakly dominated by exact matching. In particular, there are distributions μ for which the former fails to implement the perfect monopolist’s objective while the latter succeeds. The reason is that fully broad matching effectively eliminates all the information contained in consumers’ signals.

We now construct a broad-match function b^* that lies somewhere between the two extremes of exact and fully broad matching. It will be convenient to present X as a *subset* of W . This is more than a notational convenience: as in the “Mozart vs. Stravinsky” example, it is natural to assume that the name of a product type is itself a possible query. For every $w, v \in W$,

$$b^*(w | v) = \begin{cases} \sqrt{\mu(v, w)} & \text{if } v \in X \\ 0 & \text{if } v \notin X \end{cases}$$

This broad-match function has a natural interpretation. The signals in $W - X$ are “dummy keywords” with no value for bidders, because they grant no tickets to any search pool. The signals in X are “working keywords” worth bidding for. The strength of the broad match from $v = x$ to w increases with the fraction of (x, w) types among consumers.

Unlike the analysis in Section 3, here we are interested in characterizing the *entire* class of symmetric pure-strategy Nash equilibria in the game that is induced by the mechanism. Note that there are trivial equilibria in which consumers choose not to participate. The following result neglects these equilibria.

Proposition 7. *Suppose condition (8) holds for every $x, y \in X$. Then, the game induced by the broad-match auction defined by b^* has a Nash equilibrium, in which all firms of type x bid $p^*(x)$ for the signal x , where p^* is given by (9). Moreover, for generic μ , this is the unique symmetric pure-strategy equilibrium in which all consumers participate. In this equilibrium, firms surrender their entire surplus to the planner, and the induced collection of search pools is λ^* .*

Thus, b^* addresses the incentive problem due to broad matching, to the extent possible

given Proposition 4. On one hand, allocation of keywords to firms is based on competitive bidding. In equilibrium, x firms' bid-per-impression for the "working keyword" x is equal to their average conversion rate, such that their surplus is fully extracted. On the other hand, b^* ensures that the consumers' search pools are efficiently diversified, given firms' equilibrium behavior. For this, it is crucial that $b(w | x)$ is proportional to $\sqrt{\mu(x, w)}$. The relation between the broad-match auction and the direct mechanism is quite transparent: when a firm bids for x , it effectively reports that this is its type, and the broad-match function executes the role of the function q in the direct mechanism. As usual, however, the indirect mechanism generates more potential deviations, and therefore verifying equilibrium - let alone its essential uniqueness - is more intricate.

A rather restrictive feature of the broad-match auction is that firms are not allowed to bid for more than one keyword. When we relax this restriction and allow firms to bid for as many keywords as they wish, a weaker version of Proposition 7 holds: under the same conditions, the modified mechanism admits a symmetric pure-strategy Nash equilibrium that implements the perfect monopolist's objective. However, there is potentially a continuum of other equilibria, in which some firms earn positive profits.

Comment: Broad matching in sponsored ad auctions

Although our model borrows the term "broad matching" from online search, it does not mean quite the same thing, and we believe that the differences are instructive. Historically, broad matching was introduced in pursuit of two goals:⁴ (1) simplifying advertisers' bidding task, because under exact matching they need to bid for many keywords (involving typos and semantically related terms); (2) thickening markets in order to raise auction revenues. Both goals are supply-side oriented - broad matching was not designed with the vocabulary-expansion role in mind. In practice, vocabulary expansion relies on explicit attempts by the "organic" search engine (which coexists with sponsored search) to help users refine and correct their queries. In our model, goals (1) and (2) are irrelevant: if the mechanism allowed firms to bid for multiple keywords, it would be costless for them to do so; and since there are

infinitely many firms of each type, markets are already thick under exact matching. Finally, in our model the design of consumers’ search environment is entirely decentralized, hence no “organic” search engines can help consumers refine their queries. Thus, broad matching in our model is exclusively a vocabulary-expansion instrument.

As to the way broad matching interacts with the auction design, our “tickets” mechanism differs from real-life practice in two important qualitative dimensions. First, broad matching is *optional* in practice: bidders can choose to stick with exact matching. Second, in our model the winners in the auction for v are broad-matched with other queries w , whereas in practice, broad matching means that whenever a firm bids for v and selects broad matching, its bid enters auctions for other queries w . Despite these differences, the idea that broad matching mimics the vocabulary-expansion role of search intermediaries may be relevant to future work on sponsored ad auctions.

In our broad-match auction, a firm of one type does not want to kick firms of another type out of the market (by topping their bid), because this would flood the firm with “unwanted traffic” of consumers that do not generate any transactions. The possibility of receiving many impressions with infrequent conversions is a real concern in online search markets. As the following quote from Even Dar et al. (2009) shows, this concern is most profound when relying on a search engine’s broad match: “While giving more expressiveness to advertisers, this feature makes it challenging to optimize bids to maximize their returns: choosing to bid on a query as a broad match because it provides high profit results in one bidding for related queries which may yield low or even negative profits.”⁵ Of course, in our stylized model the “irrelevant traffic” takes an extreme form (infinitely many impressions), but this motive would be present in some form in a more realistic model that allows for “search fatigue” or consumer learning.

IV. Concluding Remarks

This paper addressed the following general question: under what conditions can a decentralized mechanism be efficient in helping individuals find what they want? In most of our discussion, we interpreted consumers’ signals as keyword-based queries, and this linked our model to the problem of search-engine design. In this context, our question could be rephrased as follows: would a search intermediary’s performance (measured in terms of profits or social welfare) deteriorate if it switched from a centralized matching algorithm to a decentralized mechanism?

However, our framework accommodates a wider range of environments, including ones which have yet to establish an organized marketplace for allocating firms to search pools. For example, online recommender systems give users access to search pools according to their navigation history, which serves as an imperfect signal of their current needs. In contrast to search engines, recommender systems do not purely rely on queries initiated by the web user.⁶

To see how the broad-match auction of Section 4 fits the recommendation-system interpretation, suppose that W represents a set of possible *past purchase profiles* of the consumer. In particular, we can set $W = X^K$, where K is the number of past purchase opportunities the consumer had. An element in W is a platform for “personalized advertising”, augmented by broad matching: when an advertiser pays for a particular profile of past purchases, he potentially gets access to other profiles. In this context, our question can be rephrased as follows: *suppose that a recommender system abandons its centralized recommendation algorithm in favor of a “market for sponsored recommendations”; will its performance deteriorate as a result?* Propositions 2 and 4 can be viewed as partial answers to this question.⁷

Ordered search pools

Suppose that the planner can perfectly control the order in which consumers inspect search results. Thus, he can choose which type of firm the consumer will encounter at each draw, as a function of his search history. The efficient ordering of items in a search pool

follows simple maximum likelihood: the product type displayed in the k -th position of a w consumer’s list, denoted $x_k(w)$, is the most likely to be preferred by the consumer, given that he rejected all the $k - 1$ firms whose types are $x_1(w), \dots, x_{k-1}(w)$. This ordering can be easily implemented because the firm type with the highest willingness to pay for the k -th position is the one that is most likely to transact with the consumer, which is by definition $x_k(w)$. The feature that trivializes the problem is the lack of substitutability between product types: if a consumer wants x , he will never transact with $y \neq x$. Relaxing this assumption is necessary for an interesting model with ordered search pools.

Related literature

We are not aware of precedents for our formulation of the “search design” problem. One related body of work studies mechanisms for allocating sponsored links by online search engines (e.g. Edelman et al. (2007)). Typically, the literature assumes that links have exogenous values to advertisers. Athey and Ellison (2011) explicitly model how these values are determined by consumers’ endogenous search decisions. Chen and He (2011) and Eliaz and Spiegler (2011) model explicitly the interaction between keyword and product prices (ignoring auction-theoretic considerations). This literature almost invariably assumes exact matching; two exceptions are Dhangwatnotai (2011) and Chen et al. (2014). The first study uses the “price of anarchy” framework to analyze the performance of a mechanism in which advertisers can submit a bid to multiple generalized second-price auctions at once. The second paper analyzes the worst equilibrium of a mechanism that randomly samples a keyword according to a predefined probability distribution and only runs the generalized second-price auction for this sampled keyword.

This paper is also related to the literature on intermediation in two-sided markets (see Caillaud and Jullien (2001,2003), Rochet and Tirole (2003) and Armstrong (2006)). Some works within this tradition (e.g. Hagiu and Jullien (2011)) explicitly address search platforms. One can view the consumer’s signal in our model as a (sole) platform to which he has access. In this context, broad matching can be viewed as a “*directed network of platforms*”:

a consumer who is attached to one platform is able to interact with firms that attach themselves to another platform. The papers we are aware of implicitly assume exact matching: interaction between a consumer and a firm requires that they are both attached to the same platform.⁸

Finally, in the last decade there has been much writing, both academic and popular, about the “long tail” phenomenon (see Brynjolfsson et al. (2006) or Anderson (2007)). In many industries, a large segment of the consumer population consists of numerous small taste niches. Online commerce facilitates the flourishing of the “long tail”, because it lowers barriers to entry that characterize brick-and-mortar commerce. The key friction that remains (indeed, becomes magnified) is consumers’ limited awareness of products that fit their peculiar tastes, and their limited ability to describe these tastes in order to locate relevant products on the internet. The “long tail” phenomenon means that search design can have substantial welfare implications.

REFERENCES

- Anderson, C. (2008): *The long tail: Why the future of business is selling less of more*, Hyperion Books.
- Armstrong, M. (2006): “Competition in two-sided markets,” *Rand Journal of Economics* 37(3), 668-691.
- Armstrong, M., J. Vickers and J. Zhou (2009): “Prominence and Consumer Search,” *Rand Journal of Economics* 40(2), 209-233.
- Athey, S. and G. Ellison (2011): “Position Auctions with Consumer Search,” *Quarterly Journal of Economics* 126(3), 1213-1270.
- Basu, A. H. Shioya and C. Park (2011): *Statistical Inference: The Minimum Distance Approach*, CRC Press.

- Bergmann, D. and A. Bonatti (2013): “Selling Cookies,” Cowles Foundation Discussion Paper No. 1920.
- Bhattacharyya, A. (1943): “On a measure of divergence between two statistical populations defined by their probability distributions,” *Bulletin of the Calcutta Mathematical Society* 35, 99–109.
- Brynjolfsson, E., Yu J., and M. Smith (2006): “From Niches to Riches: Anatomy of the Long Tail.” *Sloan Management Review* 47(4), 67-71.
- Caillaud, B. and B. Jullien (2001): “Competing Cybermediaries,” *European Economic Review* 45(4), 797-808.
- Caillaud, B. and B. Jullien (2003): “Chicken and Egg: Competition Among Intermediation Service Providers,” *Rand Journal of Economics* 34(2), 309-328.
- Chen, Y., C. He (2011): “Paid Placement: Advertising and Search on the Internet,” *Economic Journal* 121(556), F309-F328.
- Chen, W., D. He, T-Y. Liu, T. Qin, Y. Tao and L. Wang (2014): “Generalized Second Price Auction with Probabilistic Broad Match.” *Proceedings of the 15th ACM conference on Electronic Commerce*, 39-56.
- Cormen, T. H., Stein , C., Rivest, R. L. and C. E. Leiserson (2001): *Introduction to Algorithms*, McGraw-Hill Higher Education.
- Dhangwatnotai, P. (2011): “Multi-keyword sponsored search,” *Proceedings of the 12th ACM conference on Electronic Commerce*, 91-100.
- Edelman, B., M. Ostrovsky and M. Schwarz (2007): “Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords.” *American Economic Review* 97(1), 242-259.

- Eliasz, K. and R. Spiegler (2011): “A Simple Model of Search Engine Pricing,” *Economic Journal* 121(556), 329-339.
- Even Dar, E., V. Mirrokni, S. Muthukrishnan, Y. Mansour and U. Nadav (2009), “Bid Optimization for Broad Match Ad Auctions,” in Proceedings of the 18th International Conference on World Wide Web (pp. 231-240), ACM.
- Gomes, R. and A. Pavan (2014), “Many-to-Many Matching and Price Discrimination”, mimeo.
- Hagiu, A. and B. Jullien (2011): “Why Do Intermediaries Divert Search?,” *Rand Journal of Economics* 42(2), 337-362.
- Rochet, J-C. (1987): “A Necessary and Sufficient Condition for Rationalizability in a Quasi-linear Context,” *Journal of Mathematical Economics* 16(2), 191-200.
- Rochet, J-C. and J. Tirole (2003): “Platform Competition in Two-Sided Markets,” *Journal of the European Economic Association* 1(4), 990-1029.
- Spiegler, R. (2000): “Extracting Interaction-Created Surplus,” *Games and Economic Behavior* 30(1), 142-162.
- Theodoridis. S. and K. Koutroumbas (2008): *Pattern Recognition*, 4th Edition, Academic Press.
- Vohra, R. (2011): *Mechanism Design: A Linear Programming Approach*, Econometric Society Monograph, Cambridge University Press.
- Wolinsky, A. (1986): “True Monopolistic Competition as a Result of Imperfect Information,” *Quarterly Journal of Economics* 101(3), 493-512.

A Appendix

A. An explicit direct mechanism for Proposition 2

We construct a mechanism for each signal w (independently of all other signals), using Vohra's (2011) graphical representation of IC constraints. Consider a weighted directed graph, whose set of nodes is X , and the weight on the link $x \rightarrow y$ is $\phi(x, y) = \mu(x, w) - \sqrt{\mu(x, w)\mu(y, w)}$. Add a link from any x to itself, whose weight is $\phi(x, x) = 0$. A *path* from x to y is a sequence of nodes that begins with x and ends with y . Define the *length* of a path to be the sum of the weights on the directed links along the path. Let $\delta(x, y)$ be the distance from x to y , namely the length of the *shortest path* from x to y . Since the sum of weights along any cycle is non-negative, the distance is always well-defined and non-negative, and by definition it satisfies the triangle inequality: for any x, y, z , $\delta(x, z) \leq \delta(x, y) + \delta(y, z)$. Fix some $x^* \in X$. For any $x \in X$, define $T_x = \delta(x, x^*) - L$, where $L > 0$ is large enough to ensure that firms' participation constraints hold. By the triangle inequality, $\phi(x, y) + \delta(y, x^*) \geq \delta(x, x^*)$ for any $x, y \in X$. This implies that for any pair of distinct products x, y in X , $\phi(x, y) \geq T_x - T_y$, hence $IC(x, y)$ is satisfied.

Let us illustrate these transfers in the ‘‘Mozart vs. Stravinsky’’ example. The problem for the signal MOZ is trivial, because all consumers who submit this signal want moz . Thus, the planner can prescribe $T_{moz} = \varepsilon > 0 = T_{str}$, and if ε is small enough, no firm would have an incentive to opt out or misreport. An analogous argument holds for STR . Let us turn to the mechanism associated with CL . Note that

$$\begin{aligned}\phi(moz, str) &= \mu(moz, CL) - \sqrt{\mu(moz, CL)\mu(str, CL)} \\ \phi(str, moz) &= \mu(str, CL) - \sqrt{\mu(moz, CL)\mu(str, CL)}\end{aligned}$$

Set $T_{str} = 0$ and $T_{moz} = \phi(moz, str)$, and all participation and IC constraints will hold. ■

B. Remark 3

Denote $\mu(k | i) = \beta_{ik}$, such that (β_{ik}) is a stochastic matrix with $\sum_k \beta_{ik} = 1$ for every i .

Let

$$\delta_{ik} = \sum_h \beta_{ih} m_{hk}$$

where (m_{hk}) is a $|W| \times |W|$ bi-stochastic matrix. That it, (δ_{ik}) is a Blackwell garbling of (β_{ik}) . Fix i, j . Then,

$$\sum_k \sqrt{\delta_{ik} \delta_{jk}} = \sum_k \sqrt{\left(\sum_h \beta_{ih} m_{hk} \right) \left(\sum_h \beta_{jh} m_{hk} \right)}$$

By the Cauchy-Schwarz inequality, this expression is weakly greater than

$$\sum_k \sum_h \sqrt{\beta_{ih} m_{hk} \beta_{jh} m_{hk}} = \sum_h \sqrt{\beta_{ih} \beta_{jh}} \sum_k m_{hk} = \sum_k \sqrt{\beta_{ik} \beta_{jk}}$$

Since this inequality holds for every i, j , it follows that

$$\sum_i \sum_k \sqrt{\delta_{ik} \delta_{jk}} \geq \sum_i \sum_k \sqrt{\beta_{ik} \beta_{jk}}$$

which completes the proof. ■

C. Proof that Proposition 4 extends to mixed participation decisions

When the mechanism fully extracts firms' surplus, they are indifferent to participation. In our analysis, we restricted attention to equilibria in which all firms participate and report truthfully. The latter is w.l.o.g by the Revelation Principle. However, it is not obvious that insisting on full participation is w.l.o.g. Put differently, can we relax condition (8) by allowing firms to play a *mixed* participation strategy? Let $\alpha(x)$ be the probability that a firm of type x participates. A participating x firm has no incentive to report that its type is

y if and only if

$$\sum_{w \in W} q(x, w) \cdot \frac{\mu(x, w)}{\alpha(x)q(x, w)} - T_x \geq \sum_{w \in W} q(y, w) \cdot \frac{\mu(x, w)}{\alpha(x)q(x, w)} - T_y$$

which reduces to

$$(A1) \quad \frac{\mu(x)}{\alpha(x)} - T_x \geq \sum_{w \in W} \frac{q(y, w)}{\alpha(x)q(x, w)} \cdot \mu(x, w) - T_y$$

Because firms earn zero profits, $T_x = \mu(x)/\alpha(x)$ and $T_y = \mu(y)/\alpha(y)$. The ratio between the measures of y and x firms in the pool associated with w is $\alpha(y)q(y, w)/\alpha(x)q(x, w)$. Since the equilibrium is required to induce efficient search pools, this ratio equals $\lambda^*(y | w)/\lambda^*(x | w) = \sqrt{\mu(y, w)/\mu(x, w)}$. It follows that inequality (A1) can be written as

$$\frac{1}{\alpha(y)} \sum_{w \in W} \frac{\sqrt{\mu(y, w)}}{\sqrt{\mu(x, w)}} \cdot \mu(x, w) \leq \frac{\mu(y)}{\alpha(y)}$$

which is equivalent to condition (8). ■

D. Proof of Proposition 5

In the first-best outcome, all consumers who look for x eventually find a product they like, and they are equally shared by a measure one of x firms. Therefore, the number of transactions that an individual firm of type x expects is $\mu(x)$. The number of draws it obtains is equal to

$$\sum_w q(x, w) \sum_y \frac{\mu(y, w)}{\theta \cdot \lambda^*(y | w)} \cdot \frac{1}{\sum_y q(y, w)}$$

After we plug (2) and (4) into λ^* and q , this expression becomes

$$\frac{1}{\theta} \sum_y \sqrt{\mu(x)\mu(y)S(x, y)}$$

The ratio between the number of transactions and the number of draws is the average conversion rate experienced by x firms. Note that it is equal to $p^*(x)$. Because firms earn zero profits in the first-best outcome, the average conversion rate is equal to the amount that x firms pay on average per draw. ■

E. Proof of Proposition 7

We break the proof into three parts.

Part 1: The strategy profile induces λ^* .

Recall there is a measure one of every firm type $x \in X$. Since all x firms bid for x , the measure of tickets held by x firms in the search pool associated with any w is $b^*(w | x)$. Therefore, for every $w \in W$,

$$\lambda(x | w) = \frac{b^*(w | x)}{\sum_y b^*(w | y)} = \frac{\sqrt{\mu(x, w)}}{\sum_y \sqrt{\mu(y, w)}} = \lambda^*(x | w)$$

Part 2: The strategy profile constitutes a Nash equilibrium.

We make two preliminary observations. First, by the definition of p^* , firms earn zero profits under the assumed strategy profile. Second, the strategy profile implies that if a consumer who submits the query w draws a ticket held by an x firm, it must be a (w, x) ticket. In other words, all the tickets held by x firms in the search pool associated with w originate from the signal x . This means that if a firm of type $y \neq x$ deviates by bidding $p > p^*(x)$ for x , it will hold all the (w, x) tickets in some search pool w for which $\mu(x, w) > 0$. As a result, consumers of type (x, w) will search indefinitely without finding a product they want, and therefore the y firm will get infinitely many draws in the search pool associated with w ; the average conversion rate that the y firm will experience thanks to winning the signal x will drop to zero, and the firm's deviation will be unprofitable.

In addition, no firm has a strict incentive not to participate or to bid for a dummy signal. The former gives a payoff of zero, while the latter gives a payoff of at most zero

(when the bid is zero).

It follows that the only deviations we need to examine are of two types: (I) a firm of some type y submits $p^*(x)$ for some signal $x \neq y$; (II) a firm of some type x changes its bid for x to $p > p^*(x)$. Consider type-I deviations. Because the deviation is by a non-atomic firm that joins a measure-one set of highest bidders for x , there is a measure one of firms that hold all (w, x) tickets both before and after the deviation. Therefore, all we need to do is verify that submitting the winning bid for x generates fewer transactions for the deviating y firm than it does for the x firms that bid for x . The number of transactions for the deviating y firm is

$$(A2) \quad \sum_w \frac{b^*(w | x)}{1} \cdot \frac{\mu(y, w)}{b^*(w | y)}$$

because the firm gets $b^*(w | x)/1$ tickets to any pool associated with w . In every such pool, there is a measure $\mu(y, w)$ of consumers who are interested in y . Every such consumer eventually finds a product he wants, and this clientele is shared equally by all y firms in the pool, the measure of which is $b^*(w | y)$. By a similar calculation, the number of transactions for an x firm is

$$\sum_w \frac{b^*(w | x)}{1} \cdot \frac{\mu(x, w)}{b^*(w | x)} = \mu(x)$$

We need this to be weakly greater than (A2). Plugging the definition of b^* , we obtain

$$\mu(x) \geq \sum_w \sqrt{\mu(x, w)\mu(y, w)}$$

which is equivalent to the inequality (8). Thus, the condition that prevents type-I deviations coincides with the condition for general implementability of the planner's objective, which was assumed to hold at the outset.

Let us turn to type II deviations. Suppose that an x firm deviates by submitting a bid $p > p^*(x)$ for the signal x . The firm receives all the (w, x) tickets. However, both before

and after the deviation, all (w, x) tickets are held by x firms. Therefore, the composition of all search pools remains intact, and the number of transactions and draws per every (w, x) ticket remains unchanged. It follows that the average conversion rate associated with the signal x remains $p^*(x)$, which is below the price-per-impression that the deviating firm pays. Therefore, the firm's deviation is unprofitable.

Part 3: There are no other equilibria with full consumer participation.

We restrict attention to the generic case in which μ has full support.⁹ Let us first rule out two cases.

Case 1: The firms of at least one type x do not submit a winning bid for a working signal. This means that consumers who look for x will never be serviced. As a result, their net payoff conditional on participation is negative due to the infinite search costs, and they can profitably deviate by opting out.

Case 2: All firms submit a winning bid for some working signal, but there are two firm types x and y that bid for the same signal v . Therefore, there must be a “vacant” working signal v' that no firm bids for. Suppose that an individual x firm deviates by submitting a bid of *zero* for v' . Because μ has full support, $b^*(w | v') > 0$ for every w . Prior to the deviation, the signal v granted $b^*(w | v)$ tickets for every w , and these tickets were allocated uniformly to a population of firms of measure 2 at least. After the deviation, the deviating x firm holds all the $b^*(w | v')$ tickets granted by v' . The deviating firm thus holds infinitely many more (w, \cdot) tickets than prior to the deviation, for every w . As a result, it generates infinitely more transactions. Both before and after the deviation, all consumers eventually find a product they want, hence the number of draws that each consumer contributes is finite. Finally, thanks to the deviation, the firm pays nothing per impression. Therefore, its deviation must be profitable.

It follows that any symmetric pure-strategy equilibrium in which all consumers participate is characterized by a *permutation* $f : X \rightarrow X$, such that all x firms bid for the working signal $f(x)$. Then, all firms must earn zero profits. Assume the contrary - i.e., firms of some

type x submit a bid for $f(x)$ which lies strictly below the conversion rate they experience at $f(x)$. Then, a standard ‘‘Bertrand’’ argument applies: an individual x firm can deviate to a slightly higher bid for $f(x)$, thus acquiring an infinite number of tickets, at a slightly lower net profit per ticket than prior to the deviation. It follows that in equilibrium, the winning bid for each $y \in W^* = X$ must be equal to the conversion rate experienced by the firms that submit this bid.

Our objective is to show that the permutation f must be the *identity function*. Assume the contrary. Then, there is a collection of firm types, x_1, \dots, x_m , $m \leq |X|$, such that $f(x_k) = x_{k+1 \bmod m}$. To sustain the equilibrium, it must be the case that no firm of type $k \in \{1, \dots, m\}$ would want to deviate by mimicking the winning bid for x_k - i.e., mimicking the equilibrium behavior of firms of type $k - 1 \bmod m$. If the firm did so, it would experience the same number of draws as each of the firms whose behavior it mimics. Because all firms earn zero profits prior to the deviation, this means that all we need to check is that the deviating firm receives weakly fewer transactions.

In the putative equilibrium, each firm type x receives $\mu(x)$ transactions. The reason is that all firms of a given type x submit a winning bid for $f(x)$, and the full-support assumption implies that all consumers who look for x eventually transact with one of these firms. Let us calculate the number of transactions that a firm type $y = f(x)$ would obtain if it deviated by mimicking the behavior of x firms (i.e., mimicking their winning bid for y):

$$\sum_w \frac{b^*(w | y)}{1} \cdot \frac{\mu(y, w)}{b^*(w | f(y))} = \sum_w \frac{\mu(y, w) \sqrt{\mu(y, w)}}{\sqrt{\mu(f(y), w)}}$$

Thus, the condition that will prevent firms of type $k = 1, \dots, m$ from deviating by mimicking the behavior of $k - 1 \bmod m$ firms is

$$(A3) \quad \mu(x_{k-1 \bmod m}) \geq \sum_w \frac{\mu(x_k, w) \sqrt{\mu(x_k, w)}}{\sqrt{\mu(x_{k+1 \bmod m}, w)}}$$

By assumption, the necessary condition for the implementability of the perfect monopolist’s

objective is satisfied. In particular, for every k, j ,

$$\mu(x_k) \geq \sum_w \sqrt{\mu(x_k, w)\mu(x_j, w)}$$

Adding up a suitable selection of these inequalities, together with the inequalities (A3), we obtain the following condition:

$$2 \sum_{k=1}^m \mu(x_k) \geq \sum_{k=1}^m \sum_w \mu(x_k, w) \left[\sqrt{\frac{\mu(x_k, w)}{\mu(x_{k+1 \bmod m}, w)}} + \sqrt{\frac{\mu(x_{k+1 \bmod m}, w)}{\mu(x_k, w)}} \right]$$

For generic μ , the expression in the square brackets is strictly above 2, hence the inequality is violated. ■

Notes

*Eliaz: School of Economics, Tel-Aviv University and Economics Dept., University of Michigan, P.O.B. 39040, Ramat Aviv, Tel-Aviv 69978, Israel, kfire@post.tau.ac.il. Spiegler: School of Economics, Tel-Aviv University and Economics Dept., University College London, P.O.B. 39040, Ramat Aviv, Tel-Aviv 69978, rani@post.tau.ac.il. This is a radically revised version of a paper that appeared under different titles: "The Market for Keywords" and "Broad Matching and the Market for Search Platforms". We gratefully acknowledge financial support from ISF grant, no. 1153/13. We thank Yair Antler, Yossi Azar, Alex Frug, Stephen Hansen, Sergiu Hart, Sebastien Lahaie, David Pennock, Justin Rao, Neil Thakral, the editor and referees of this journal, and numerous seminar participants for helpful comments. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

¹We assume this symmetry purely for notational simplicity; it is easy to adapt our analysis to the case of profit margins that vary across firm types.

²More generally, each firm is assigned a probability distribution over *subsets* of search pools. However, because firms' payoff is separable across search pools, it is w.l.o.g. to associate with each report a probability distribution over *individual* search pools.

³See Basu, Shioya and Park (2011) and Theodoris and Koutroumbas (2008). A related concept is the *Hellinger distance* between distributions, given by $H^2(x, y) = 1 - \sqrt{S(x, y)}$.

⁴This historical account is informed by discussions with David Pennock (a pioneer of sponsored-link auction formats, which employed broad matching) and Justin Rao. We are grateful to them for sharing their knowledge and insights, and apologize for any imprecision in this rendering.

⁵In a similar spirit, an account manager at Bing Ads writes in a blog post titled "The People vs. Bing Broad Match" that one of the most common complaints levied against Bing broad match is that "it drives irrelevant traffic". See <http://www.ppchero.com/the-people-vs-bing-broad-match/>

⁶Netflix automatically displays movie recommendations for its subscribers on its homepage; when a consumer buys a particular product on Amazon, the checkout screen displays recommended products, even though the consumer was not actively searching for these products; and when a researcher views a scholarly article on ScienceDirect, the side panel displays links to other recommended articles.

⁷For a model of pricing of targeted ads, see Bergemann and Bonatti (2013).

⁸Some papers in the literature on two-sided markets have taken a mechanism-design approach. Spiegler (2000) examines contract design by an intermediary who aims to match agents who could interact elsewhere, and extract their joint surplus. Gomes and Pavan (2014) study mechanisms for implementing efficient many-to-many matching when agents are privately informed about their "vertical" characteristics.

⁹This genericity requirement can be weakened, at the cost of having a considerably longer proof.