The Benefits of Search Costs in Multiagent Exploration (Extended Abstract)

David Sarne
Bar-Ilan University
david.sarne@gmail.com

Yonatan Aumann
Bar-Ilan University
yaumann@gmail.com

Categories and Subject Descriptors

J.4 [**Social and Behavioral Sciences**]: Economics

General Terms

Algorithms, Economics

Keywords

Sequential Search, Two-Sided Search, Social Welfare

1. INTRODUCTION

Humans and software agents alike spend considerable time and effort in searching. Search enables finding the things that better fit and agent's goals. But search can also be a costly process. Search costs can either come in the form of direct monetary payments, or in the form of time and resources spent. In general, the searcher must balance between the benefits provided by longer and broader search, on the one hand, and the associated increased cost, on the other.

In economic literature search costs are often referred to as "environment friction" or "market *inefficiency*" and associated with reduced market performance [1]. Indeed, in the presence of search costs a rational player will not aim to find the best option, but rather settle for the "good enough", beyond which the marginal cost of searching exceeds the marginal benefit of continuing the search. Thus, search costs promote sub-optimal results (or so it would seem). As such, the traditional wisdom is that when designing a MAS environment, search costs should be avoided or reduced to a minimum. Taking eCommerce as an example, most researchers see a great benefit in the ability of eMarketplaces to lower the buyers' cost to obtain information (e.g. about the price features) from multiple sellers, as well as the sellers' reduced costs to communicate their information [1]. The lowered search cost is associated in this case with increased economic efficiency and enable new markets to emerge. Similarly, many systems have been introduced in which central mechanisms or mediators are used in order to supply the agents complete information concerning market opportunities, eliminating the need to engage in costly search.

In this paper we show that, not withstanding the above, search costs – "friction", if and when applied appropriately, can also be beneficial, and actually *increase* both the expected aggregate welfare, and the expected utility of each individual agent. This holds even if the proceeds from the search costs are discarded and do not directly benefit anyone in the system, as we assume throughout. Thus, artificially applied search costs can be used as a mechanism to improve market efficiency. We show this for one-sided search settings using standard models from *search theory.* Similar results for two-sided search settings are available however omitted for space considerations.

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), 4-8 June 2012, Valencia, Spain.

2. ONE-SIDED SEARCH

The Model.

We employ the fixed-sample-size search model [2], wherein each searcher executes a single search round in which it obtains a large set of opportunities simultaneously, and chooses the one associated with the highest utility. Consider an environment with m homogeneous servers, and N homogeneous agents requesting service from these servers. The agents are assigned a random order. Each agent, in turn, can request to query any number of servers. The queries are assigned to available servers. Each server can address one query in each time step. Since the queries may be executed in parallel, agents need to determine the number of queries they request in advance. Once the queries are executed, the agent obtains the results and leaves the system.

Each reply provides the agent with some non-negative utility. The utility, x, obtained by agent A_j from the reply of server i is randomly drawn from a distribution D_U characterized by a probability distribution function (p.d.f.) $f(x)$, and cumulative distribution function (c.d.f.) $F(x)$. For simplicity we assume that all servers and all agents are homogeneous, and thus share the same functions $f(x)$ and $F(x)$. The overall utility obtained by the agent from the set of all replies is the maximum among the utilities provided by the individual replies from the different servers. We assume that the future is discounted by a factor of δ (which is common to all agents).

Agents are assumed to be self-interested, and thus aim to maximize their own expected utility. Thus, if there is no cost for querying a server, all agents will request to query all servers. This, however, means that serving each agent takes more resources, and thus delays serving other agents. Since the future is discounted, agents further down the line actually end up losing from this delay more than they gain from accessing more serves. We show that by introducing a cost for each query, we can drive agents to perform less queries, and increase the expected utility.

No Search Costs.

For a search cost of c, let S_c be the expected aggregate utility with a search cost of c. For any $k = 1, \dots, m$, let $E_k = E(\max\{x_1, \dots, x_k : x \in D_U\})$ be the expected maximum of k independent draws from the utility distribution D_U. Then, if there are no search costs, ignoring discounting, each agent would obtain expected utility of E_m. However, each agent occupies all the m available servers. Thus, the i-th agent is only serviced at time i. Thus, taking discounting into account, the overall expected utility of the i-th agent is $E_m\delta^i$. The aggregate expected utility, summed up over all agents, is thus:

$$S_0 = \sum_{i=1}^{N} E_m\delta^i = \delta E_m \frac{1-\delta^N}{1-\delta} \underset{N\to\infty}{\longrightarrow} \frac{\delta E_m}{1-\delta} \quad (1)$$

With Search Costs.

Now, assume that we introduce a cost c for each query. Then, the rational choice for an agent is to query $k \le m$ servers such that the expected marginal utility of querying the k-th server (rather than $k-1$ servers) is at least c, but the marginal utility of querying the $k+1$ server is less than c. Thus, each agent will choose to query k servers such that $k = \arg\max_k\{E_k - E_{k-1} \ge c\}$. Conversely, the minimum cost that will guarantee querying exactly k servers is: $c_k = E_{k+1} - E_k$.

With a search cost of c_k, ignoring discounting, the expected utility of each agent is: $U_k = E_k - k\cdot c_k = (k+1)E_k - kE_{k+1}$. At any one time step, m/k agents can be served in parallel (assuming k divides m). Thus, the i-th agent is served at time $\left\lfloor \frac{i}{m/k} \right\rfloor$. Thus, taking into account discounting, the expected utility of the i-th agent is $U_k\delta^{\left\lfloor \frac{i}{m/k} \right\rfloor}$. Thus, the total expected utility, summed up over all agents, is:

$$S_{c_k} = \sum_{i=1}^{\left\lfloor \frac{N}{m/k} \right\rfloor} \frac{m}{k} U_k\delta^i = \frac{m}{k}\delta U_k \frac{1-\delta^{\left\lfloor \frac{N}{m/k} \right\rfloor}}{1-\delta} \underset{N\to\infty}{\longrightarrow} \frac{m}{k}\frac{\delta U_k}{1-\delta} \quad (2)$$

Advantageous Search Costs - Aggregate Utility.

THEOREM 1. *For any non-degenerate distribution D_U on non-negative utilities and any discounting factor $\delta < 1$, there exist m_0 and N_0 such that for any $m \ge m_0$ and $N \ge N_0$, there exists a c such that introducing a search cost of c for each query increases the expected aggregate utility. This holds even if the proceeds of the search costs are discarded and do not benefit anyone.*

PROOF. Set $\vec{S_0} = \frac{\delta E_m}{1-\delta}$ and $\vec{S_{c_k}} = \frac{m}{k}\frac{\delta U_k}{1-\delta}$. By (1) and (2) we have that $S_0 \underset{N\to\infty}{\longrightarrow} \vec{S_0}$ and $S_{c_k} \underset{N\to\infty}{\longrightarrow} \vec{S_{c_k}}$. Suppose that $\vec{S_{c_k}} > \vec{S_0}$. Set $\epsilon = \vec{S_{c_k}} - \vec{S_0}$. Then, there exists an N_0 such that for any $N \ge N_0$, $S_0 < \vec{S_0} + \epsilon/2$ and $S_{c_k} \ge \vec{S_{c_k}} - \epsilon/2$. Thus, for $N \ge N_0$ we have that $S_{c_k} > S_0$, i.e. introducing a search cost of c_k increases aggregate utility.

It thus remains to prove that $\vec{S_{c_k}} > \vec{S_0}$. We show that this holds for any k, provided that m is sufficiently large. Indeed, $\vec{S_{c_k}} = \frac{m}{k}\frac{\delta U_k}{1-\delta} > \frac{\delta E_m}{1-\delta} = \vec{S_0}$ iff

$$\frac{U_k}{k} > \frac{E_m}{m} \quad (3)$$

The left hand side of (3) is independent of m, while the right hand side approaches 0 as m grows [2]. Thus, provided that U_k is positive, (3) necessarily holds for m sufficiently large.

We show that U_k is positive for any k. Denote $f_k(x)$ the p.d.f. of the maximum of k independent samples from D_U, and let $F_k(x)$ be the associated c.d.f. Then, $F_k(x) = (F(x))^k$ and $f_k(x) = (F_k(x))' = k(F(x))^{k-1}f(x)$. By definition, $E_k = \int_0^\infty yf_k(y)dy$. Thus,

$$U_k = (k+1)E_k - kE_{k+1} =$$
$$(k+1)\int_0^\infty f_k(y)ydy - k\int_0^\infty f_{k+1}(y)ydy =$$
$$k(k+1)\int_0^\infty (F(y))^{k-1}(1-F(y))f(y)ydy > 0$$

The last inequality is due to the fact that all elements of the integral are non-negative, and assuming that the distribution is non-degenerate (i.e. is not concentrated all in one value) at least one element is strictly positive. □

Individual Utility.

COROLLARY 1. *For any non-degenerate distribution D_U on non-negative utilities and any discounting factor $\delta < 1$, if agents are assigned a random order then there exist m_0 and N_0 such that for any $m \ge m_0$ and $N \ge N_0$, there exists a c such that introducing a search cost of c for each query increases the expected utility of each player. This holds even if the proceeds of the search costs are discarded and do not benefit anyone.*

PROOF. Considering a specific player, for any position i, the probability that the player is i-th in the order is $1/N$. Thus, if there are no search costs then the expected utility of any player is:

$$\sum_{i=1}^{N}\frac{1}{N}E_m\delta^i = \frac{1}{N}S_0 \quad (4)$$

Similarly, the expected utility of the player with a search cost of c_k, is

$$\sum_{i=1}^{\left\lfloor \frac{N}{m/k} \right\rfloor}\frac{1}{N}\frac{m}{k}U_k\delta^i = \frac{1}{N}S_{c_k} \quad (5)$$

Thus, the theorem follows by the exact same reasoning as that in the proof of Theorem 1. □

3. CONCLUSIONS

The implication of these results to market designers is that the effects of search costs should be carefully analyzed in each case, and not assumed to be universally detrimental. Rather, there are cases when it may be beneficial to deliberately introduce artificial search costs. When search costs are already part of the system, there is no general answer for whether or not decreasing these costs will improve the system's performance. In some settings, an increase rather than a decrease can actually contribute to improving expected utility. In other cases, a decrease in search costs can contribute to improving expected utility, but decreasing the costs beyond a certain point can result with the opposite effect. The analysis methodology given in this paper can facilitate the calculation of the right search cost to which the market designer should strive.

Acknowledgments
This work was supported in part by ISF/BSF grants 1401/09 and 2008-404.

4. REFERENCES

[1] Y. Bakos. Reducing buyer search costs: Implications for electronic marketplaces. *Mgmt. Sci.*, 42:1676–92, 1997.

[2] G. Stigler. The economics of information. *Journal of Political Economy*, 69(3):213–225, 1961.