

Sequential Decision Making in Parallel Two-Sided Economic Search

David Sarne
School of Engineering and Applied Sciences
Harvard University
Cambridge MA 02138 USA

Teijo Arponen
Institute of Mathematics
Helsinki University of Technology
SF-02015 TKK, Finland

ABSTRACT

This paper presents a two-sided economic search model in which agents are searching for beneficial pairwise partnerships. In each search stage, each of the agents is randomly matched with several other agents in parallel, and makes a decision whether to accept a potential partnership with one of them. The distinguishing feature of the proposed model is that the agents are not restricted to maintaining a synchronized (instantaneous) decision protocol and can sequentially accept and reject partnerships within the same search stage. We analyze the dynamics which drive the agents' strategies towards a stable equilibrium in the new model and show that the proposed search strategy weakly dominates the one currently in use for the two-sided parallel economic search model. By identifying several unique characteristics of the equilibrium we manage to efficiently bound the strategy space that needs to be explored by the agents and propose an efficient means for extracting the distributed equilibrium strategies in common environments.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*Intelligent agents*

General Terms

Algorithms, Economics

Keywords

Two-Sided Search, Matching

1. INTRODUCTION

A two-sided economic search is a distributed mechanism for forming agents' pairwise partnerships [5].¹ On every stage of the process, each of the agents is randomly matched with another agent

¹Notice that the concept of "search" here is very different from the classical definition of "search" in AI. While AI search is an active process in which an agent finds a sequence of actions that will bring it from the initial state to a goal state, economic search refers to the identification of the best agent to commit to a partnership with.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'07 May 14–18 2007, Honolulu, Hawai'i, USA.
Copyright 2007 IFAAMAS .

and the two interact bilaterally in order to learn the benefit encapsulated in a partnership between them. The interaction does not involve bargaining thus each agent merely needs to choose between accepting or rejecting the partnership with the other agent. A typical market where this kind of two-sided search takes place is the marriage market [22]. Recent literature suggests various software agent-based applications where a two-sided distributed (i.e., with no centralized matching mechanisms) search takes place. An important class of such applications includes secondary markets for exchanging unexploited resources. An exchange mechanism is used in those cases where selling these resources is not the core business of the organization or when the overhead for selling them makes it non-beneficial. For example, through a two-sided search, agents, representing different service providers, can exchange unused bandwidth [21] and communication satellites can transfer communication with a greater geographical coverage. Two-sided agents-based search can also be found in applications of buyers and sellers in eMarkets and peer-to-peer applications. The two-sided nature of the search suggests that a partnership between a pair of agents is formed only if it is mutually accepted. By forming a partnership the agents gain an immediate utility and terminate their search. When resuming the search, on the other hand, a more suitable partner might be found however some resources will need to be exchanged for maintaining the search process.

In this paper we focus on a specific class of two-sided search matching problems, in which the performance of the partnership applies to both parties, i.e., both gain an equal utility [13]. The equal utility scenario is usually applicable in domains where the partners gain from the synergy between them. For example, consider tennis players that seek partners when playing doubles (or a canoe's paddler looking for a partner to practice with). Here the players are being rewarded completely based on the team's (rather than the individual) performance. Other examples are the scenario where students need to form pairs for working together on an assignment, for which both partners share the same grade, and the scenario where two buyer agents interested in similar or interchangeable products join forces to buy a product together, taking advantage of discount for quantity (i.e. each of them enjoys the same reduced price). In all these applications, any two agents can form a partnership and the performance of any given partnership depends on the skills or the characteristics of its members. Furthermore, the equal utility scenario can also hold whenever there is an option for side-payments and the partnership's overall utility is equally split among the two agents forming it [22].

While the two-sided search literature offers comprehensive equilibrium analysis for various models, it assumes that the agents' search is conducted in a purely sequential manner: each agent locates and interacts with one other agent in its environment at a time

[5, 22]. Nevertheless, when the search is assigned to autonomous software agents a better search strategy can be used. Here an agent can take advantage of its unique inherent filtering and information processing capabilities and its ability to efficiently (in comparison to people) maintain concurrent interactions with several other agents at each stage of its search. Such use of parallel interactions in search is favorable whenever the average cost² per interaction with another agent, when interacting in parallel with a batch of other agents, is smaller than the cost of maintaining one interaction at a time (i.e., advantage to size). For example, the analysis of the costs associated with evaluating potential partnerships between service providers reveals both fixed and variable components when using the parallel search, thus the average cost per interaction decreases as the number of parallel interactions increases [21].

Despite the advantages identified for parallel interactions in adjacent domains (e.g., in one-sided economic search [7, 16]), a first attempt for modeling a repeated pairwise matching process in which agents are capable of maintaining interaction with several other agents at a time was introduced only recently [21]. However, the agents in that seminal model are required to synchronize their decision making process. Thus each agent, upon reviewing the opportunities available in a specific search stage, has to notify all other agents of its decision whether to commit to a partnership (at most with one of them) or reject the partnership (with the rest of them). This inherent restriction imposes a significant limitation on the agents' strategic behavior.

In our model, the agents are free to notify the other agents of their decisions in an asynchronous manner. The asynchronous approach allows the agents to re-evaluate their strategy, based on each new response they receive from the agents they interact with. This leads to a sequential decision making process by which each agent, upon sending a commit message to one of the other agents, delays its decision concerning a commitment or rejection of all other potential partnerships until receiving a response from that agent (i.e., the agent still maintains parallel interactions in each search stage, except that its decision making process at the end of the stage is sequential rather than instantaneous). The new model is a much more realistic pairwise model and, as we show in the analysis section, is always preferred by any single agents participating in the process. In the absence of other economic two-sided parallel search models, we use the model that relies on an instantaneous (synchronous) decision making process [21] (denoted I-DM throughout the rest of the paper) as a benchmark for evaluating the usefulness of our proposed sequential (asynchronous) decision making strategy (denoted S-DM).

The main contributions of this paper are threefold: First, we formally model and analyze a two-sided search process in which the agents have no temporal decision making constraints concerning the rejection of or commitment to potential partnerships they encounter in parallel (the S-DM model). This model is a general search model which can be applied in various (not necessarily software agents-based) domains. Second, we prove that the agents' S-DM strategy weakly dominates the I-DM strategy, thus every agent has an incentive to deviate to the S-DM strategy when all other agents are using the I-DM strategy. Finally, by using an innovative recursive presentation of the acceptance probabilities of different potential partnerships, we identify unique characteristics of the equilibrium strategies in the new model. These are used for supplying an appropriate computational means that facilitates the calculation of the agents' equilibrium strategy. This latter contribution is

²The term "costs" refers to resources the agent needs to consume for maintaining its search, such as: self advertisement, locating other agents, communicating with them and processing their offers.

of special importance since the transition to the asynchronous mode adds inherent complexity to the model (mainly because now each agent needs to evaluate the probabilities of having each other agent being rejected or accepted by each of the other agents it interacts with, in a multi-stage sequential process). We manage to extract the agents' new equilibrium strategies without increasing the computational complexity in comparison to the I-DM model. Throughout the paper we demonstrate the different properties of the new model and compare it with the I-DM model using an artificial synthetic environment.

In the following section we formally present the S-DM model. An equilibrium analysis and computational means for finding the equilibrium strategy are provided in Section 3. In Section 4 we review related MAS and economic search theory literature. We conclude with a discussion and suggest directions for future research in Section 5.

2. MODEL AND ANALYSIS

We consider an environment populated with an infinite number of self-interested fully rational agents of different types³. Any agent A_i can form a partnership with any other agent A_j in the environment, associated with an immediate perceived utility $U(A_i, A_j)$ for both agents. As in many other partnership formation models (see [5, 21]) we assume that the value of $U(x, y)$ (where x and y are any two agents in the environment) is randomly drawn from a continuous population characterized with a probability distribution function (p.d.f.) $f(U)$ and a cumulative distribution function (c.d.f.) $F(U)$, ($0 \leq U < \infty$). The agents are assumed to be acquainted with the utility distribution function $f(x)$, however they cannot tell a-priori what utility can be gained by a partnership with any specific agent in their environment. Therefore, the only way by which an agent A_i can learn the value of a partnership with another agent A_j , $U(A_i, A_j)$, is by interacting with agent A_j . Since each agent in two-sided search models has no prior information concerning any of the other agents in its environment, it initiates interactions (i.e., search) with other agents randomly. The nature of the two-sided search application suggests that the agents are satisfied with having a single partner, thus once a partnership is formed the two agents forming it terminate their search process and leave the environment.

The agents are not limited to interacting with a single potential partner agent at a time, but rather can select to interact with several other agents in parallel. We define a search round/stage as the interval in which the agent interacts with several agents in parallel and learns the utility of forming a partnership with each of them. Based on the learned values, the agent needs to decide whether to commit or reject each of the potential partnerships available to it. Commitment is achieved by sending a commit message to the appropriate agent and an agent cannot commit to more than one potential partnership simultaneously. Declining a partnership is achieved by sending a reject message. The communication between the agents is assumed to be asynchronous and each agent can delay its decision, concerning any given potential partnership, as necessary.⁴ If two agents A_i and A_j mutually commit to a partnership between

³The infinite number of agents assumption is common in two-sided search models (see [5, 22, 21]). In many domains (e.g., eCommerce) this derives from the high entrance and leave rates, thus the probability of running into the same agent in a random match is negligible.

⁴Notice that the asynchronous procedure does not eliminate the inherent structure of the search. The search is still based on stages/rounds where on each search round the agent interacts with several other agents, except that now the agent can delay its decision making process (within each search round) as necessary.

them, then the partnership is formed and both agents gain the immediate utility $U(A_i, A_j)$ associated with it. If an agent does not form a partnership in a given search stage, it continues to its next search stage and interacts with more agents in a similar manner.

Given the option for asynchronous decision making, each individual agent, A_i , follows the following procedure:

```

1: loop
2:   Set  $N$  (number of parallel interactions for next search round)
3:   Locate randomly a set  $A = \{A_1, \dots, A_N\}$  of agents to interact with
4:   Evaluate the set of utilities  $\{U(A_i, A_1), \dots, U(A_i, A_N)\}$ 
5:   Set  $A^* = \{A_j | A_j \in A \text{ and } U(A_i, A_j) > U(\text{resume})\}$ 
6:   Send a reject message to each agent in the set  $\{A \setminus A^*\}$ 
7:   while ( $A^* \neq \emptyset$ ) do
8:     Send a commit message to  $A_j = \text{argmax}_{A_i \in A^*} U(A_i, A_j)$ 
9:     Remove  $A_j$  from  $A^*$ 
10:    Wait for  $A_j$ 's decision
11:    if ( $A_j$  responded "commit") then
12:      Send "reject" messages to the remaining agents in  $A^*$ 
13:      Terminate search
14:    end if
15:  end while
16: end loop

```

where $U(\text{resume})$ denotes the expected utility of continuing the search (in the following paragraphs we show that $U(\text{resume})$ is fixed throughout the search and derives from the agent's strategy). In the above algorithm, any agent A_i first identifies the set A^* of other agents it is willing to accept out of those reviewed in the current search stage and sends a reject message to the rest. Then it sends a commit message to the agent $A_j \in A^*$ that is associated with the partnership yielding the highest utility. If a reject message was received from agent A_j then this agent is removed from A^* and a new commit message is sent according to the same criteria. The process continues until either: (a) the set A^* becomes empty, in which case the agent initiates another search stage; or (b) a dual commitment is obtained, in which case the agent sends reject messages to the remaining agents in A^* . The method differs from the one used in the I-DM model in the way it handles the commitment messages: in the I-DM model, after evaluating the set of utilities (step 4), the agent merely sends instantaneously a commit message to the agent associated with the greatest utility and a reject message to all the other agents it interacted with (as a replacement to steps 5-15 in the above procedure). Our proposed S-DM model is much more intuitive as it allows an agent to "hold" and possibly exploit relatively beneficial opportunities even if its first priority partnership is rejected by the other agent. In the I-DM model, on the other hand, since reject messages are sent alongside the commit message, simultaneously, a reject message from the agent associated with the "best" partnership enforces a new search round.

Notice that the two-sided search mechanism above aligns with most other two-sided search mechanisms in a sense that it is based on "random matching" (i.e., in each search round the agent encounters a random sample of agents). While the maintenance of the random matching infrastructure is an interesting research question, it is beyond the scope of this paper. Notwithstanding, we do wish to emphasize that given the large number of agents in the environment and the fact that in MAS the turnover rate is quite substantial due to the open nature of the environment (and the interoperability between environments). Therefore, the probability of ending up interacting with the same agent more than once, when initiating a random interaction, is practically negligible.

THEOREM 1. *The S-DM agent's decision making process: (a) is the optimal one (maximizes the utility) for any individual agent*

in the environment; and (b) guarantees a zero deadlock probability for any given agent in the environment.

Proof:

(a) The method is optimal since it cannot be changed in a way that produces a better utility for the agent. Since bargaining is not applicable here (benefits are non-divisible) then the agent's strategy is limited to accepting or rejecting offers. The decision of rejecting a partnership in step 6 is based only on the immediate utility that can be gained from this partnership in comparison to the expected utility of resuming the search (i.e., moving on to the next search stage) and is not affected by the willingness of the other agents to commit or reject a partnership with A_i . As for partnerships that yield a utility greater than the expected utility of resuming the search (i.e., the partnerships with agents from the set A^*), the agent always prefers to delay its decision concerning partnerships of this type until receiving all notifications concerning potential partnerships that are associated with a greater immediate utility. The delay never results with a loss of opportunity since the other agent's decision concerning this opportunity is not affected by agent A_i 's willingness to commit or reject this opportunity (but rather by the other agent's estimation of its expected utility if resuming the search and the rejection messages it receives for more beneficial potential partnerships). Finally, the agent cannot benefit from delaying a commit message to the agent associated with the highest utility in A^* , thus will always send it a commit message.

(b) We first prove the following lemma that states that the probability of having two partnering opportunities associated with an identical utility is zero.

LEMMA 2.1. *When f is a continuous distribution function, then*

$$\lim_{y \rightarrow x} \left(\left[\int_{z=x}^y f(z) dz \right]^2 \right) = 0.$$

Proof: since f is continuous and the interval between x and y is finite, by the intermediate value theorem (found in most calculus texts) there exists a c between x and y that

$$\int_{z=x}^y f(z) dz = f(c)(y-x)$$

(intuitively, a rectangle with the base from $z = x$ to $z = y$ and height $= f(c)$ has the same area as the integral on the left hand side.). Therefore

$$\left[\int_{z=x}^y f(z) dz \right]^2 = |f(c)|^2 |y-x|^2$$

When $y \rightarrow x$, $f(c)$ stays bounded due to continuity of f , moreover $\lim_{y \rightarrow x} f(c) = f(x)$, hence

$$\lim_{y \rightarrow x} \left(\left[\int_{z=x}^y f(z) dz \right]^2 \right) = f(x)^2 \lim_{y \rightarrow x} |y-x|^2 = 0. \quad \square$$

An immediate derivative from the above lemma is that no tie-breaking procedures are required and an agent in a waiting state is always waiting for a reply from the single agent that is associated with the highest utility among the agents in the set A^* (i.e., no other agent in the set A^* is associated with an equal utility). A deadlock can be formed only if we can create a cyclic sequence of agents in which any agent is waiting for a reply from the subsequent agent in the sequence. However, in our method any agent A_i will be waiting for a reply from another agent A_j , to which it sent a commit message, only if: (1) any agent $A_k \in A$, associated with a utility $U(A_i, A_k) > U(A_i, A_j)$, has already rejected the partnership with agent A_i ; and (2) agent A_j itself is waiting for a reply from agent A_l where $U(A_l, A_j) > U(A_j, A_i)$. Therefore, if we have a sequence of waiting agents then the utility associated with partnerships between any two subsequent agents in the sequence must increase along the sequence. If the sequence is cyclic, then we have a

pattern of the form: $U(A_i, A_l) > U(A_l, A_j) > U(A_j, A_i)$. Since $U(A_i, A_l) > U(A_j, A_i)$, agent A_i can be waiting for agent A_j only if it has already been rejected by A_l (see (1) above). However, if agent A_l has rejected agent A_i then it has also rejected agent A_j . Therefore, agent A_j cannot be waiting for agent A_l to make a decision. The same logic can be applied to any longer sequence. \square

The search activity is assumed to be costly [11, 1, 16] in a way that any agent needs to consume some of its resources in order to locate other agents to interact with, and for maintaining the interactions themselves. We assume utilities and costs are additive and that the agents are trying to maximize their overall utility, defined as the utility from the partnership formed minus the aggregated search costs along the search process. The agent's cost of interacting with N other agents (in parallel) is given by the function $c(N)$. The search cost structure is principally a parameter of the environment and thus shared by all agents.

An agent's strategy $S(A') \rightarrow \{\text{commit } A_j \in A', \text{ reject } A'' \subset A', N\}$ defines for any given set of partnership opportunities, A' , what is the subset of opportunities that should be immediately declined, to which agent to send a commit message (if no pending notification from another agent is expected) or the number of new interactions to initiate (N). Since the search process is two-sided, our goal is to find an equilibrium set of strategies for the agents.

2.1 Strategy Structure

Recall that each agent declines partnerships based on (a) the partnerships' immediate utility in comparison to the agent's expected utility from resuming search; and (b) achieving a mutual commitment (thus declining pending partnerships that were not rejected in (a)). Therefore an agent's strategy can be represented by a pair (N^t, x^t) where N^t is the number of agents with whom it chooses to interact in search stage t and x^t is its reservation value⁵ (a threshold) for accepting/rejecting the resulting N potential partnerships. The subset A^* , thus, will include all partnership opportunities of search stage t that are associated with a utility equal to or greater than x^t . The reservation value x^t is actually the expected utility for resuming the search at time t (i.e., $U(\text{resume})$). The agent will always prefer committing to an opportunity greater than the expected utility of resuming the search and will always prefer to resume the search otherwise.

Since the agents are not limited by a decision horizon, and their search process does not imply any new information about the market structure (e.g., about the utility distribution of future partnership opportunities), their strategy is stationary - an agent will not accept an opportunity it has rejected beforehand (i.e., $x^1 = x^2 = \dots = x$) and will use the same sample size, $N^1 = N^2 = \dots = N$, along its search.

2.2 Calculating Acceptance Probabilities

The transition from instantaneous decision making process to a sequential one introduces several new difficulties in extracting the agents' strategies. Now, in order to estimate the probability of being accepted by any of the other agents, the agent needs to recursively model, while setting its strategy, the probabilities of rejections other agents might face from other agents they interact with. In the following paragraphs we introduce several complementary definitions and notations, facilitating the formal introduction of the acceptance probabilities. Consider an agent A_i , using a strategy (N, x_N) while operating in an environment where all other agents

⁵Notice the reservation value used here is different from a reservation price concept (that is usually used as buyers' private evaluation). The use of reservation-value based strategies is common in economic search models [21, 17].

are using a strategy (k, x_k) . The probability that agent A_i will receive a commitment message from agent A_j it interacted with depends on the utility associated with the potential partnership between them, x . This probability, denoted by $G_k(x)$ can be calculated as:⁶

$$G_k(x) = \begin{cases} \left(1 - \int_{y=x}^{\infty} f(y)G_k(y)dy\right)^{k-1} & \text{if } x \geq x_k \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

The case where $x < x_k$ above is trivial: none of the other agents will accept agent A_i if the utility in such a partnership is smaller than their reservation value x_k . However even when the partnership's utility is greater or equal to x_k , commitment is not guaranteed. In the latter scenario, a commitment message from agent A_j will be received only if agent A_j has been rejected by all other agents in its set A^* that were associated with a utility greater than the utility of a partnership with agent A_i .

The unique solution to the recursive Equation 1 is:

$$G_k(x) = \begin{cases} \left(1 + (k-2) \int_{y=x}^{\infty} f(y)dy\right)^{\frac{1-k}{k-2}}, & k > 2, x \geq x_k, \\ \exp\left(-\int_{y=x}^{\infty} f(y)dy\right), & k = 2, x \geq x_k, \\ 1, & k = 1, x \geq x_k \\ 0, & x < x_k. \end{cases} \quad (2)$$

Notice that as expected, a partnership opportunity that yields the maximum mutual utility is necessarily accepted by both agents, i.e., $\lim_{x \rightarrow \infty} G_k(x) = 1$. On the other hand, when the utility associated with a potential partnership opportunity is zero ($x = 0$) the acceptance probability is non-negligible:

$$\lim_{x \rightarrow 0} G_k(x) = (k-1)^{\frac{1-k}{k-2}} \quad (3)$$

This non-intuitive result derives from the fact that there is still a non-negligible probability that the other agent is rejected by all other agents it interacts with.

2.3 Setting the Agents' Strategies

Using the function $G_k(x)$, we can now formulate and explore the agents' expected utility when using their search strategies. Consider again an agent A_i that is using a sample of size N while all other agents are using a strategy (k, x_k) . We denote by $R_N(x)$ the probability that the maximum utility that agent A_i can be guaranteed when interacting with N agents (i.e., the highest utility to which a commit message will be received) is at most x . This can be calculated as the probability that none of N agents send agent A_i a commit message for a partnership associated with a utility greater than x :

$$R_N(x) = \left(1 - \int_{\max(x, x_k)}^{\infty} f(y)G_k(y)dy\right)^N \quad (4)$$

Notice that $R_N(x)$ is in fact a cumulative distribution function, satisfying: $\lim_{x \rightarrow \infty} R_N(x) = 1$ and $dR_N(x)/dx > 0$ (the function never gets a zero value simply because there is always a positive probability that none of the agents commit at all to a partnership with agent A_i). Therefore, the derivative of the function $R_N(x)$, denoted $r_N(x)$, is in fact the probability distribution function of the maximum utility that can be guaranteed for agent A_i when sampling N other agents:

$$r_N(x) = \frac{dR_N(x)}{dx} = \begin{cases} Nf(x)G_k(x)^{\frac{N+k-2}{k-1}}, & x \geq x_k \\ 0, & x < x_k \end{cases} \quad (5)$$

⁶The use of the recursive Equation 1 is enabled since we assume that the number of agents is infinite (thus the probability of having an overlap between the interacting agents and the affect of such overlap on the probabilities we calculate become insignificant).

This function $r_N(x)$ is essential for calculating $V_N(x_N)$, the expected utility of agent A_i when using a strategy (N, x_N) , given the strategy (k, x_k) used by the other agents:

$$V_N(x_N) = \int_{y=\max(x_N, x_k)}^{\infty} y r_N(y) dy + \left(1 - \int_{y=\max(x_N, x_k)}^{\infty} r_N(y) dy\right) V_N(x_N) - c(N) \quad (6)$$

The right hand side of the above equation represents the expected utility of agent A_i from taking an additional search stage. The first term represents the expected utility from mutual commitment scenarios, whereas the second term is the expected utility associated with resuming the search (which equals $V_N(x_N)$ since nothing has changed for the agent). Using simple mathematical manipulations and substituting $r_N(x)$, Equation 6 transforms into:

$$V_N(x) = \frac{\int_{y=\max(x, x_k)}^{\infty} y N f(y) G_k(y)^{\frac{N+k-2}{k-1}} dy - c(N)}{\int_{y=\max(x, x_k)}^{\infty} N f(y) G_k(y)^{\frac{N+k-2}{k-1}} dy} \quad (7)$$

and further simplified into:

$$V_N(x) = \max(x, x_k) + \frac{\int_{\max(x, x_k)}^{\infty} (1 - G_k(y)^{\frac{N}{k-1}}) dy - c(N)}{1 - G_k(\max(x, x_k))^{\frac{N}{k-1}}} \quad (8)$$

Equation 8, allows us to prove some important characteristics of the model as summarized in the following Theorem 2.

THEOREM 2. *When other agents use strategy (k, x_k) :*

(a) *An agent's expected utility function, $V_N(x_N)$, when using a strategy (N, x) , is quasi concave in x with a unique maximum, obtained for the value x_N satisfying:*

$$V_N(x_N) = x_N \quad (9)$$

(b) *The value x_N satisfies:*

$$c(N) = (\max(x_N, x_k) - x_N) \left(1 - G_k(x_k)^{\frac{N}{k-1}}\right) + \int_{\max(x_N, x_k)}^{\infty} (1 - G_k(y)^{\frac{N}{k-1}}) dy \quad (10)$$

The proof is obtained by deriving $V_N(x_N)$ in Equation 8 and setting it to zero. After applying further mathematical manipulations we obtain (9) and (10).

Both parts of Theorem 2 can be used as an efficient means for extracting the optimal reservation value x_N of an agent, given the strategies of the other agents in the environment and the number of parallel interactions it uses. Furthermore, in the case of complex distribution functions where extracting x_N from Equation 10 is not immediate, a simple algorithm (principally based on binary search) can be constructed for calculating the agent's optimal reservation value (which equals its expected utility, according to 9), with a complexity $O(\log(\frac{\hat{x}}{\rho}))$, where ρ is the required precision level for x_N and \hat{x} is the solution to: $\int_{y=\hat{x}}^{\infty} y N f(y) F(y)^{N-1} dy = c(N)$.

Having the ability to calculate x_N , we can now prove the following Proposition 2.1.

PROPOSITION 2.1. *An agent operating in an environment where all agents are using a strategy according to the instantaneous parallel search equilibrium (i.e., according to the I-DM model [21]) can only benefit from deviating to the proposed S-DM strategy.*

Sketch of proof: For the I-DM model the following holds [21]:

$$c(N) = \frac{N}{2N-1} \int_{y=x_N}^{\infty} (1 - F(y))^{2N-1} dy \quad (11)$$

We apply the methodology used above in this subsection for constructing the expected utility of the agent using the S-DM strategy as a function of its reservation value, assuming all other agents are using the I-DM search strategy. This results with an optimal reservation value for the agent using S-DM, satisfying:

$$c(N) = \int_{y=x_N^{S-DM}}^{\infty} \left(1 - \left(1 - \frac{1}{N} + \frac{F(y)^N}{N}\right)^N\right) dy \quad (12)$$

Finally, we prove that the integrand in Equation 11 is smaller than the integrand in Equation 12. Given the fact that both terms equal $c(N)$, we obtain $x_N^{S-DM} > x_N^{I-DM}$ and consequently (according to Theorem 2) a similar relationship in terms of expected utilities. \square

Figure 1 illustrates the superiority of the proposed search strategy S-DM, as well as the expected utility function's characteristics (as reflected in Theorem 2). For comparative reasons we use the same synthetic environment that was used for the I-DM model [21]. Here the utilities are assumed to be drawn from a uniform distribution function and the cost function was taken to be $c(N) = 0.05 + 0.005N$. The agent is using $N = 3$ while other agents are using $k = 25$ and $x_k = 0.2$. The different curves depict the expected utility of the agent as a function of the reservation value, x , that it uses, when: (a) all agents are using the I-DM strategy (marked as I-DM); (b) the agent is using the S-DM strategy while the other agents are using the I-DM strategy (marked as I-DM/S-DM); and (c) all agents are using the S-DM strategy (marked as S-DM). As expected, according to Equation 8 and Theorem 2, the agent's expected utility remains constant until its reservation value exceeds x_k . Then, it reaches a global maximum when the reservation value satisfies $V_N(x) = x$. From the graph we can see that the agent always has an incentive to deviate from the I-DM strategy to S-DM strategy (as was proven in Proposition 2.1).

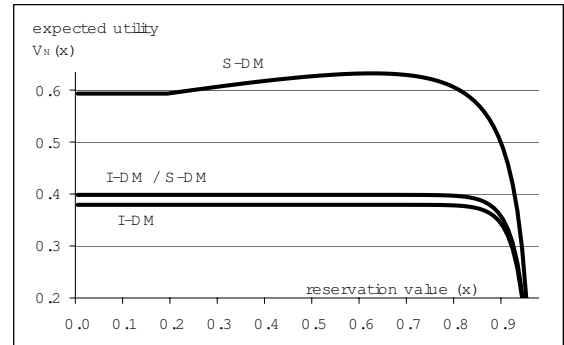


Figure 1: The expected utility as a function of the reservation value used by the agent

3. EQUILIBRIUM DYNAMICS

Since all agents are subject to similar search costs, and their perceived utilities are drawn from the same distribution function, they all share the same strategy in equilibrium. A multi-equilibria scenario may occur, however as we discuss in the following paragraphs since all agents share the same preferences/priorities (unlike, for example, in the famous "battle of the sexes" scenario) we can always identify which equilibrium strategy will be used.

Notice that if all agents are using the same sample size, N , then the value x_N resulting from solving Equation 10 by substituting $k = N$ and $x_k = x_N$ is a stable reservation value (i.e., none of the agents can benefit from changing just the value of x_N).

An equilibrium strategy (N, x_N) can be found by identifying an N value for which no single agent has an incentive to use a different number of parallel interactions, k (and the new optimal reservation

value that is associated with k according to Equation 10). While this implies an infinite solution space, we can always bound it using Equations 8 and 10. Within the framework of this paper, we demonstrate such a bounding methodology for the common case where $c(N)$ is linear⁷ or convex, by using the following Theorem 3.

THEOREM 3. *When $c(N)$ is linear (or convex), then: (a) When all other agents sample k potential partners over a search round, if an agent's expected utility of sampling $k + 1$ potential partners, $V_{k+1}(x_{k+1})$, is smaller than $V_k(x_k)$, then the expected utility when sampling N potential partners, $V_N(x_N)$, where $N > k + 1$, is also smaller than $V_k(x_k)$. (b) Similarly, when all other agents sample k potential partners over a search round, if an agent's expected utility of using $k - 1$ potential partners, $V_{k-1}(x_{k-1})$, is smaller than the expected utility when using k potential partners, $V_k(x_k)$, then the expected utility when using N potential partners, where $N < k - 1$, is also smaller than $V_k(x_k)$.*

Proof: Let us use the notation c_i for $c(i)$. Since $V_k(x_k) = x_k \quad \forall k$ (according to Equation 9), the claims are: (a) if $x_{k+1} < x_k$ then $x_N < x_k$ for all $N \geq k + 1$, and (b) if $x_{k-1} < x_k$ then $x_N < x_k$ for all $N \leq k - 1$.

(a) We start by proving that if $x_{k+1} < x_k$ then $x_{k+2} < x_k$. Assume otherwise, i.e., $x_{k+1} < x_k$ and $x_{k+2} > x_k$. Therefore, according to Equation 10, the following holds:

$$0 < c_{k+2} - 2c_{k+1} + c_k < \int_{x_{k+2}}^{\infty} (1 - G_k(y))^{\frac{k+2}{k-1}} dy - 2 \int_{x_k}^{\infty} (1 - G_k(y))^{\frac{k+1}{k-1}} dy + \int_{x_k}^{\infty} (1 - G_k(y))^{\frac{k}{k-1}} dy$$

where the transition to inequality is valid since $c(i)$ is convex. Since the assumption in this proof is that $x_{k+2} > x_k$ then the above can be transformed into:

$$\int_{x_k}^{\infty} (2G_k(y)^{\frac{k+1}{k-1}} - G_k(y)^{\frac{k+2}{k-1}} - G_k(y)^{\frac{k}{k-1}}) dy > 0 \quad (13)$$

Now notice that the integrated term is actually $-G_k(y)^{\frac{k}{k-1}} (1 - G_k(y)^{\frac{1}{k-1}})^2$ which is obviously negative, contradicting the initial assumption, thus if $x_{k+1} < x_k$ then necessarily $x_{k+2} < x_k$.

Now we need to prove the same for any x_{k+j} . We will prove this in two steps: first, if $x_{k+i} < x_k$ then $x_{k+2i} < x_k$. Second, if $x_{k+i} < x_k$ and $x_{k+i+1} < x_k$, then $x_{k+2i+1} < x_k$. Together these constitute the necessary induction arguments to prove the case (a). We start with the even case, using a similar methodology: Assume otherwise, i.e., $x_{k+l} < x_k \quad \forall l = 1, \dots, j - 1$ and $x_{k+2i} > x_k$. According to Equation 10, and the fact that $c(i)$ is convex, the following holds:

$$\int_{x_k}^{\infty} (2G_k(y)^{\frac{k+i}{k-1}} - G_k(y)^{\frac{k+2i}{k-1}} - G_k(y)^{\frac{k}{k-1}}) dy > 0 \quad (14)$$

And again the integrand is actually $-G_k(y)^{\frac{k}{k-1}} (1 - G_k(y)^{\frac{i}{k-1}})^2$ which is obviously negative, contradicting the initial assumption, thus $x_{k+2i} < x_k$.

As for the odd case, we use Equation 10 once for $k + i + 1$ parallel interactions and once for $k + 2i + 1$. From the convexity of c_i , we obtain: $c_{k+2i+1} - c_{k+i} - c_{k+i+1} + c_k > 0$, thus:

$$\int_{x_k}^{\infty} (G_k(y)^{\frac{k+i}{k-1}} + G_k(y)^{\frac{k+i+1}{k-1}} - G_k(y)^{\frac{k+2i+1}{k-1}} - G_k(y)^{\frac{k}{k-1}}) dy > 0 \quad (15)$$

⁷A linear cost function is mostly common in agent-based two-sided search applications, since often the cost function can be divided into fixed costs (e.g. operating the agent per time unit) and variable costs (i.e., cost of processing a single interaction's data).

This time the integrated term in Equation 15 can be re-written as $G_k(y)^{\frac{k}{k-1}} (1 - G_k(y)^{\frac{i}{k-1}}) (G_k(y)^{\frac{i+1}{k-1}} - 1)$ which is obviously negative, contradicting the initial assumption, thus $x_{k+i+1} < x_k$.

Now using induction one can prove that if $x_{k+1} < x_k$ then $x_{k+i} < x_k$. This concludes part (a) of the proof.

The proof for part (b) of the theorem is obtained in a similar manner. In this case: $c_k - 2c_{k-i} + c_{k-2i} > 0$ and $c_k - c_{k-i-1} - c_{k-i} + c_{k-2i-1} > 0$. \square

The above theorem supplies us with a powerful tool for eliminating non-equilibrium N values. It suggests that we can check the stability of a sample size N and the appropriate reservation value x_N simply by calculating the optimal reservation values of a single agent when deviating towards using samples of sizes $N - 1$ and $N + 1$ (keeping the other agents with strategy (N, x_N)). If both the appropriate reservation values associated with the two latter sample sizes are smaller than x_N then according to Theorems 3 the same holds when deviating to any other sample size k . The above process can be further simplified by using $V_{N+1}(x_N) > x_N$ and $V_{N-1}(x_N) > x_N$ as the two elimination rules. This derives from Theorem 3 and the properties of the function $V_N(x)$ found in Theorem 2.

Notice that a multi-equilibria scenario may occur, however can easily be resolved. If several strategies satisfy the stability condition defined above, then the agents will always prefer the one associated with the highest expected utility. Therefore an algorithm that goes over the different N values and checks them according to the rules above can be applied, assuming that we can bound the interval for searching the equilibrium N . The following Theorem 4 suggests such an upper bound.

THEOREM 4. *An upper bound for the equilibrium number of partners to be considered over a search round is the solution of the equation:*

$$A(N) = c(N) \quad (16)$$

provided $A(N - 1) > c(N - 1)$, where we denote,

$$A(N) := \int_{y=0}^{\infty} yNf(y)G_k(y)^{\frac{N+k-2}{k-1}} dy.$$

Proof: We denote:

$$A(N, x) = \int_{y=x}^{\infty} yNf(y)G_k(y)^{\frac{N+k-2}{k-1}} dy$$

so that $A(N) = A(N, 0)$. From Equation 7:

$$V_N(x) = \frac{A(N, x) - c(N)}{N \int_x^{\infty} f(y)G_k(y)^b dy} = \frac{A(N, x) - c(N)}{\text{positive}},$$

Clearly $A(N) \geq A(N, x) \forall x$ since the integrand is positive. Hence if $A(N) - c(N) < 0$, then $A(N, x) - c(N) < 0 \forall x$ and $V_N(x) < 0 \forall x$. Next we prove that if $A(N) - c(N)$ gets negative, it stays negative.

Recalling that for any $g(y)$:

$$\frac{d}{dN} (g(y)^{b(N)}) = g(y)^{b(N)} \log(g(y)) \frac{db}{dN}$$

we get:

$$A''(N) = \frac{-1}{(k-1)^2} \int_0^{\infty} G_k(y)^{\frac{N}{k-1}} (\log G_k(y))^2 dy$$

which is always negative, since the integrand is nonnegative. Therefore $A(N)$ is concave. Since $c(N)$ is convex, $-c(N)$ is concave, and a sum of concave functions is concave, we obtain that

$A(N) - c(N)$ is concave. This guarantees that once the concave expression $A(N) - c(N)$ shifts from a positive value to a negative one (with the increase in N), it cannot become positive again. Therefore, having N^* such that $A(N^*) = c(N^*)$, and $A(N^{**}) > c(N^{**})$ for some $N^{**} < N^*$, is an upper bound for N , i.e., $V_N(x) < 0 \quad \forall N \geq N^*$. The condition we specify for N^{**} is merely for ensuring that V_N is switching from a positive value to a negative one (and not vice versa) and is trivial to implement. \square

Given the existence of the upper bound, we can design an algorithm for finding the equilibrium strategy (if one exists). The algorithm extracts the upper bound, \bar{N} , for the equilibrium number of parallel interactions according to Theorem 4. Out of the set of values satisfying the stability condition defined above, the algorithm chooses the one associated with the highest reservation value according to Equation 10. This is the equilibrium associated with the highest expected utility to all agents according to Theorem 2.




Figure 2: The incentive to deviate from strategy (N, x_N)

The process is illustrated in Figure 2 for an artificial environment where partnerships' utilities are associated with a uniform distribution. The cost function used is $c(N) = 0.2 + 0.02N$. The graph depicts a single agent's expected utility when all other agents are using N parallel interactions (on the horizontal axis) and the appropriate reservation value x_N (calculated according to Equation 10). The different curves depict the expected utility of the agent when it uses a strategy: (a) (N, x_N) similar to the other agents (marked as $V_N(x_N)$); (b) $(N + 1, x_N)$ (marked as $V_{N+1}(x_N)$); and (c) $(N - 1, x_N)$ (marked as $V_{N-1}(x_N)$). According to the discussion following Theorem 3, a stable equilibrium satisfies: $V_N(x_N) > \max\{V_{N+1}(x_N), V_{N-1}(x_N)\}$. The strategy satisfying the latter condition in our example is $(9, 0.437)$.

4. RELATED WORK

The two-sided economic search for partnerships in AI literature is a sub-domain of coalition formation⁸. While coalition formation models usually consider general coalition-sizes [24], the partnership formation model (often referred as matchmaking) considers environments where agents have a benefit only when forming a partnership and this benefit can not be improved by extending the partnership to more than two agents [12, 23] (e.g., in the case of buyers and sellers or peer-to-peer applications). As in the general

⁸The use of the term "partnership" in this context refers to the agreement between two individual agents to cooperate in a pre-defined manner. For example, in the buyer-seller application a partnership is defined as an agreed transaction between the two-parties [9].

coalition formation case, agents have the incentive to form partnerships when they are incapable of executing a task by their own or when the partnership can improve their individual utilities [14]. Various centralized matching mechanisms can be found in the literature [6, 2, 8]. However, in many MAS environments, in the absence of any reliable central matching mechanism, the matching process is completely distributed.

While the search in agent-based environments is well recognized to be costly [11, 21, 1], most of the proposed coalition formation mechanisms assume that an agent can scan as many partnership opportunities in its environment as needed or have access to central matchers or middle agents [6]. The incorporation of costly search in this context is quite rare [21] and to the best of our knowledge, a distributed two-sided search for partners model similar to the S-DM model has not been studied to date.

Classical economic search theory ([15, 17], and references therein) widely addresses the problem of a searcher operating in a costly environment, seeking to maximize his long term utility. In these models, classified as one-sided search, the focus is on establishing the optimal strategies for the searcher, assuming no mutual search activities (i.e., no influence on the environment). Here the sequential search procedure is often applied, allowing the searcher to investigate a single [15] or multiple [7, 19] opportunities at a time. While the latter method is proven to be beneficial for the searcher, it was never used in the "two-sided" search models that followed (where dual search activities are modeled) [22, 5, 18]. Therefore, in these models, the equilibrium strategies are always developed based on the assumption that the agents interact with others sequentially (i.e., with one agent at a time). A first attempt to integrate the parallel search into a two-sided search model is given in [21], as detailed in the introduction section.

Several of the two-sided search essences can be found in the strategic theory of bargaining [3] - both coalition formation and matching can be represented as a sequential bargaining game [4] in which payoffs are defined as a function of the coalition structure and can be divided according to a fixed or negotiated division rule. Nevertheless, in the sequential bargaining literature, most emphasis is put on specifying the details of the sequential negotiating process over the division of the utility (or cost) jointly owned by parties or the strategy the coalition needs to adopt [20, 4]. The models presented in this area do not associate the coalition formation process with search costs, which is the essence of the analysis that economic search theory aims to supply. Furthermore, even in repeated pairwise bargaining [10] models the agents are always limited to initiating a single bargaining interaction at a time.

5. DISCUSSION AND CONCLUSIONS

The phenomenal growth evidenced in recent years in the number of software agent-based applications, alongside the continuous improvement in agents' processing and communication capabilities, suggest various incentives for agents to improve their search performance by applying advanced search strategies such as parallel search. The multiple-interactions technique is known to be beneficial for agents both in one-sided and two-sided economic search [7, 16, 21], since it allows the agents to decrease their average cost of learning about potential partnerships and their values. In this paper we propose a new parallel two-sided search mechanism that differs from the existing one in a sense that it allows the agents to delay their decision making process concerning the acceptance and rejection of potential partnerships as necessary. This, in comparison to the existing instantaneous model [21] which force each agent to make a simultaneous decision concerning each of the potential partnerships revealed to it during the current search stage.

As discussed throughout the paper, the new method is much more intuitive to the agent than the existing model - an agent will always prefer to keep all options available. Furthermore, as we prove in the former sections, an agent's transition to the new search method always results with a better utility.

As we prove in Section 2, in spite of the transition to a sequential decision making, deadlocks never occur in the proposed method as long as all agents use the proposed strategies. Since our analysis is equilibrium-based, a deviation from the proposed strategies is not beneficial. Similarly, we show that a deviation of a single agent (back) to the instantaneous decision making strategy is not beneficial. The only problem that may arise in the transition from an instantaneous to sequential decision making is when an agent fails (technically) to function (endlessly delaying the notification to the agents it interacted with). While equilibrium analysis normally do not consider malfunction as a legitimate strategy, we do wish to emphasize that the "malfunctioning agent" problem can be resolved by using a simple timeout for receiving responses and skipping this agent in the sequential decision process if the timeout is exceeded.

Our analysis covers all aspects of the new two-sided search technique, from individual strategy construction throughout the dynamics that lead to stability (equilibrium). The difficulty in the extraction of the agents' equilibrium strategies in the new model derives from the need to recursively model, while setting an agent's strategy, the rejection other agents might face from other agents they interact with. This complexity (that does not exist in former models) is resolved by the introduction of the recursive function $G_k(x)$ in Section 2. Using the different theorems and propositions we prove, we proffer efficient tools for calculating the agents' equilibrium strategies. Our capabilities to produce an upper bound for the number of parallel interactions used in equilibrium (Theorem 4) and to quickly identify (and eliminate) non-equilibrium strategies (Theorem 3) resolves the problem of the computational complexity associated with having to deal with a theoretically infinite strategy space.

While the analysis we present is given in the context of software agents, the model we suggest is general, and can be applied to any two-sided economic search environment where the searchers can search in parallel. In particular, in addition to weakly dominating the instantaneous decision making model (as we prove in the analysis section) the proposed method weakly dominates the purely sequential two-sided search model (where each agent interacts with only one other agent at a time) [5]. This derives from the fact that the proposed method is a generalization of the latter (i.e., in the worst case scenario, the agent can interact with one other agent at a time "in parallel").

Naturally the attempt to integrate "search theory" techniques into day-to-day applications brings up the applicability question. Justification and legitimacy considerations for this integration were discussed in the wide literature we refer to throughout the paper. The current paper is not focused on re-arguing applicability, but rather on the improvement of the the core two-sided search model. We see great importance in future research that will combine bargaining as part of the interaction process. We believe such research can result in many rich variants of our two-sided search model.

6. REFERENCES

- [1] Y. Bakos. Reducing buyer search costs: Implications for electronic marketplaces. *Management Science*, 42(12):1676–1692, June 1997.
- [2] G. Becker. A theory of marriage. *Journal of Political Economy*, 81:813–846, 1973.
- [3] K. Binmore, M. Osborne, and A. Rubinstein. Non-cooperative models of bargaining. In *Handbook of Game Theory with Economic Applications*, pages 180–220. Elsevier, New York, 1992.
- [4] F. Bloch. Sequential formation of coalitions in games with externalities and fixed payoff division. *Games and Economic Behavior*, 14(1):90–123, 1996.
- [5] K. Burdett and R. Wright. Two-sided search with nontransferable utility. *Review of Economic Dynamics*, 1:220–245, 1998.
- [6] K. Decker, K. Sycara, and M. Williamson. Middle-agents for the internet. In *Proc. of IJCAI*, pages 578–583, 1997.
- [7] S. Gal, M. Landsberger, and B. Levyskon. A compound strategy for search in the labor market. *Int. Economic Review*, 22(3):597–608, 1981.
- [8] D. Gale and L. Shapley. College admissions and the stability of marriage. *American Math. Monthly*, 69:9–15, 1962.
- [9] M. Hadad and S. Kraus. Sharedplans in electronic commerce. In M. Klusch, editor, *Intelligent Information Agents*, pages 204–231. Springer Publisher, 1999.
- [10] M. Jackson and T. Palfrey. Efficiency and voluntary implementation in markets with repeated pairwise bargaining. *Econometrica*, 66(6):1353–1388, 1998.
- [11] J. Kephart and A. Greenwald. Shopbot economics. *JAAMAS*, 5(3):255–287, 2002.
- [12] M. Klusch. Agent-mediated trading: Intelligent agents and e-business. *J. on Data and Knowledge Engineering*, 36(3), 2001.
- [13] S. Kraus, O. Shehory, and G. Taase. Coalition formation with uncertain heterogeneous information. In *Proc. of AAMAS '03*, pages 1–8, 2003.
- [14] K. Lermann and O. Shehory. Coalition formation for large scale electronic markets. In *Proc. of ICMAS'2000*, pages 216–222, Boston, 2000.
- [15] S. A. Lippman and J. J. McCall. The economics of job search: A survey. *Economic Inquiry*, 14:155–189, 1976.
- [16] E. Manisterski, D. Sarne, and S. Kraus. Integrating parallel interactions into cooperative search. In *AAMAS*, pages 257–264, 2006.
- [17] J. McMillan and M. Rothschild. Search. In R. Aumann and S. Hart, editors, *Handbook of Game Theory with Economic Applications*, pages 905–927. 1994.
- [18] J. M. McNamara and E. J. Collins. The job search problem as an employer-candidate game. *Journal of Applied Probability*, 27(4):815–827, 1990.
- [19] P. Morgan. Search and optimal sample size. *Review of Economic Studies*, 50(4):659–675, 1983.
- [20] A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–109, 1982.
- [21] D. Sarne and S. Kraus. Agents strategies for the dual parallel search in partnership formation applications. In *Proc. of AMEC2004, LNCS 3435*, pages 158 – 172, 2004.
- [22] R. Shimer and L. Smith. Assortative matching and search. *Econometrica*, 68(2):343–370, 2000.
- [23] K. Sycara, S. Widoff, M. Klusch, and J. Lu. Larks: Dynamic matchmaking among heterogeneous software agents in cyberspace. *JAAMAS*, 5:173–203, 2002.
- [24] N. Tsvetovat, K. Sycara, Y. Chen, and J. Ying. Customer coalitions in electronic markets. In *Proc. of AMEC2000*, pages 121–138, 2000.