Managing Parallel Inquiries in Agents' Two-Sided Search

David Sarne
School of Engineering and Applied Sciences,
Harvard University, Cambridge MA 02138 USA
sarned@eecs.harvard.edu

Sarit Kraus
Department of Computer Science,
Bar-Ilan University, Ramat Gan, 52900 Israel
sarit@cs.biu.ac.il

August 29, 2007

Abstract

In this paper we address the problem of agents engaged in a distributed costly two-sided search for pairwise partnerships in Multi-Agent Systems (MAS). While traditional two-sided search mechanisms are based on a purely sequential search of all searchers, our mechanism integrates an ability of some of the agents to maintain several search efforts in parallel at each search stage. We show that in many environments the transition to the new mechanism is inevitable since the adoption of the parallel-interactions based search suggests a greater utility for the searching agents. By exploring the appropriate model equations, we present the new dynamics that drive the equilibrium when using such a mechanism in MAS environments. Complementary algorithms are offered, based on the unique equilibria characteristics found, for facilitating the extraction of the agents' strategies. The analysis methodology used supplies a comprehensive solution to a self contained model, and also offers a great value for future work concerning distributed two-sided mechanisms for MAS. Towards the end of the paper we review two of these important models that can benefit from the proposed analysis.

Keywords: Multi-Agent Systems, Autonomous Agents, Equilibrium Analysis, Matching

1 Introduction

In this paper we consider the problem of agents engaged in a distributed costly two-sided search for partners [8] in Multi-Agent Systems (MAS). The problem is often classified as a matching problem, since the agents' goal is to form pairwise partnerships. In this problem each agent is satisfied with only one partner and gains no utility from extending (upon finding a partner) its partnership further or from operating on its own. The matching problem is a unique variant of the general coalition formation model and its main incentive is similar to the one which drives coalitions of agents: throughout partnering, the agents can operate more effectively and coordinate their activities [39], thus increase the participants' benefits [7].

Various centralized matching mechanisms can be found in literature [10, 3, 14]. However, in many MAS environments, in the absence of any reliable central matching mechanism, the matching process is completely distributed. In a distributed matching model the agents need to search for partnering opportunities. The agents learn about new partnering opportunities (and the benefits encapsulated in them) through bilateral interactions. The search process is considered two-sided when all agents in the environment engage in search. Thus a partnership eventually formed is the result of the combined search activities of both sides of the interaction (i.e., the agents forming it). Similarly, 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 them.

This concept of two-sided search for forming partnerships can be found in many traditional economical applications such as the marriage market [40] and the job market [27]. It can also be found in many MAS applications [22], e.g., buyer and seller agents operating in electronic marketplaces and peer-to-peer distributed applications.[1] An important class of such applications includes secondary markets for exchanging unexploited resources. An exchange mechanism is used in those cases where selling the 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 [37] and communication satellites can transfer communication with a greater geographical coverage. In all these applications an agent can gain a utility only if it eventually partners with another agent. However, once a partnership is formed, adding more agents as partners does not produce any additional benefit.

The main idea of this paper is that a distributed two-sided search in MAS environments should take into consideration agents' capability to use parallel (simultaneous) interactions with other agents. This is in comparison to the traditional models found in the two-sided search literature [8, 40] where the agents' search is conducted in a purely sequential manner: each agent locates and interacts with one agent in its environment at a time.[2] Autonomous computer agents have unique inherent filtering and information processing capabilities and, most important, the ability to efficiently (in comparison to people) maintain concurrent interactions with several other agents at each given time [15, 2, 21]. This way an agent can make a decision at each stage of its search based on interactions with several other agents (instead of one). Such use of parallel interactions in search is favorable especially when the search is costly, as explained

[1] The 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 [16].

[2] As we report in the related work section, the use of parallel search was suggested for problems of a single searcher [29]. However, these were merely optimization problems that did not concern equilibrium dynamics.

in the following paragraphs. The transition into using this new search technique results in new dynamics in MAS environments and consequently new equilibrium structures. Specifically, we consider environments where the agents are associated with two possible types (i.e., a buyer and a seller) and only agents of one of the types use the parallel interaction search method. While this model is self contained and associated with specific applications (as illustrated in the following paragraphs) it can also be used for understanding the dynamics formed in models where agents of all types use parallel search. The various aspects concerning the transition to the latter model are discussed towards the end of the paper.

The motivation for using the parallel search technique is mainly the existence of search costs. These costs are a common inherent part of MAS environments where there is no central source that can supply full immediate reliable information on the environment and the different opportunities that can be found in it. The costs reflect the resources (not necessarily monetary) that need to be invested/consumed by an agent to perform its search activities (e.g., the cost associated with the interaction and negotiation between agents, locating other agents, analyzing and comparing offers, decision making, self-advertisement and the cost of maintaining the agent in an idle state until finding a partnership). The introduction of search costs leads to a more realistic description of MAS environments. Many authors have argued that recent advances in communication technologies reduce search costs and other environmental inefficiencies in MAS environments [6]. However the general agreement is that these cannot be ignored completely [2] and should be integrated into the agents' expected utility computation process, given their specific search strategy. Others have argued that the search cost for locating an opportunity is insignificant compared to the utility encapsulated in most opportunities. Nevertheless the growing interoperability between different systems and environments in the internet age, followed by a phenomenal increase in the number and complexity of opportunities available, makes the overall cost of acquiring such information an important parameter that needs to be considered when forming the agents' strategies [9, 21, 36].

Given the search cost, the key issue for each agent engaged in a distributed two-sided search is to determine the set of agents with whom it is willing to form a partnership. By forming a partnership the agent gains an immediate utility and terminates its search. On the other hand, if the search is resumed, a more suitable partner might be found, but some resources will need to be consumed for maintaining the extended search process. The advantage of parallel search within this context is mainly in reducing the average cost of an interaction with an agent of the other type. This reduction is achieved whenever the cost of interacting with a batch of agents is smaller than the overall cost of interacting with each of them sequentially. For example, in [37] 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.

The main contributions of this paper are threefold: First, we formally model and analyze a distributed two-sided search model in which agents of a specific type can use parallel interactions in their search for a pairwise partnership. This model is a general search model which can be applied in various (not necessarily computer-agent based) domains. As detailed in the next sections, the adoption of the new search technique creates new dynamics (mutual strategies adjustments) that drive the agents' strategies towards a stable equilibrium (different from the purely sequential models). Second, we show that our mechanism is a generalization of the traditional purely sequential two-sided search model, thus the agent's utility will never decrease when using our proposed mechanism. As we illustrate, the agents have a strong incentive to deviate from the sequential search strategy to the new strategy in many environments. The task of extracting the agents' new equilibrium strategies adds some computational complexity. This is mainly because now the agents are not limited to strategies in which they only need to decide what agents to accept, but rather can influence the intensity of their search by deciding on the number of agents with which they will interact in parallel. Therefore, finally, by using the unique characteristics identified for the agents strategies and for the equilibrium structure we supply appropriate algorithms that facilitate the calculation of the agents' equilibrium strategy.

As a framework for our analysis we use a legacy two-sided agents' search application - the buyer-seller search in an electronic marketplace. Issues concerning the design of agent-mediated electronic trading systems involve finding solutions for a diverse set of interaction problems. These range from behavioral to organizational issues and also encompass complex computational, information and system level challenges [32, 18, 15, 41]. Our specific focus in ecommerce is on the C2C (Consumer-to-Consumer) segment, where a transaction is always associated with two consumers. In C2C we usually find non-repeated transactions and thus no a-priori information concerning specific buyers and sellers is collected prior to the buyer's "need identification" stage of the process.[3] Therefore each agent needs to explore the market for opportunities to buy or sell according to its owner's personal preferences and requirements. Notice that the seller agents in C2C marketplaces usually have a single item (or a limited quantity) they wish to sell on an irregular basis. Therefore, even though the seller agents do not search proactively for buyer agents, they act in a selective manner in order to maximize their total utility.[4] Also, in the C2C application, typically, buyer agents are the ones capable of searching in parallel since they are the ones that approach the seller type agents.

Additional MAS-based two-sided search applications in which only agents of a specific type can search in parallel can be found in the task/resource allocation domain. Consider, for example a scenario of self-interested servers that offer their computational services to agents in an open MAS environment. Since each server has a limited set of resources it may prefer to reject "non-profitable" jobs (assuming each server charges a (different) fixed price for execution, thus the profit is a function of the time it takes to execute a specific task). The agents, on the other hand, try to minimize the payment to the servers for processing their jobs. Since the jobs vary in their characteristics and the servers highly vary in their configuration and load, there is no general correlation between a job and the time it takes to execute it on any specific server. Here, only agents can search in parallel (interact with servers to learn the cost of executing the job), since the servers do not know the identity of the incoming agents/jobs prior to the time they approach them. Another MAS-based interesting application that has been recently introduced in this domain is the one where experts offer their services, for example as in kasamba.com.[5] Since those who seek for service are not listed on the web-site (only the service providers are listed), they are the only ones who have the capability of proactively addressing service providers (and

[3] This is the stage in the Consumer Buying Behavior (CBB) model, in which the consumer realizes her need for the product [15].

[4] Notice that the concept of "search" in this case 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, in the buyer-seller domain search refers to the identification of the best agent to sell to or buy from. In general, the concept here refers to the determination of with whom to commit to a partnership.

[5] People use Kasamba when they need immediate professional guidance, or help on a project. Nowadays, one can find more than 30,000 professionals offering their services to people and hundreds of thousands of visitors seeking service arriving to the web-site every month.

thus search in parallel). The service seekers' objective is to maximize their utility which is a function of the money they spend on the service, the time they spend looking for the service provider (and the interaction with experts along the way) and the quality of the service/solution they obtain. The service providers, on the other hand, have a resource constraint (i.e., an expert has limited time throughout the day in which she can engage in providing the service) thus an expert's utility function depends on the money she obtains, the amount of time she needs to invest in this service, the deadline set for this service (and the way it aligns with other jobs it received (or expects to receive)), etc.

In the following section we formally describe the model and detail the agents' interaction mechanism. We derive the mechanism's innovative characteristics from an analytic investigation of the agents' expected benefit equations, presented in section 3. We expand the analysis in that section to a complete equilibrium analysis by exploring the dynamics that drive the agents' strategies given the strategies set by the other agents in their environment. Throughout the analysis section we use a specific synthetic artificial environment to illustrate the unique characteristics of the model and the equilibrium dynamics. In addition, we supply efficient algorithms for extracting the agents' strategies that can be used by market makers and agent designers. In section 4 we discuss two important future extensions of the model and draw the guidelines for their analysis. A review of relevant work both in multi-agent systems and in the related area of economical search theory is given in section 5. We conclude with a discussion concerning the generality and applicability aspects of the suggested model (section 6).

2 The Model

We begin by describing the model and its assumptions, and provide the appropriate notations which are used later in our analysis sections.

2.1 Model Description

We consider an environment populated with numerous self-interested agents of two types, where each agent is interested in forming a partnership with a single agent of the opposite type. For illustrative purposes, in our model description we adopt the eCommerce domain terminology. Therefore, the two types can be seen as buyer and seller agents, respectively, residing in a C2C marketplace environment. Each agent is interested in buying or selling a specific item as described below. Being in a dynamic environment with a high rate of entrance and exit of agents, the agents have no a-priori valuation concerning the utility that can be obtained by partnering with specific agents of the other type in the environment. In order to learn this information, an agent needs to interact with the other agents (according to given conventions defined by ontology and a language). This is in the absence of a centralized trusted mediator with global immediate knowledge that can direct the agents into partnerships they all accept. Having no a-priori information, each agent interacts randomly with others and if the buyer agent's preferences and requirements for the product attributes and functionalities are met, then a possible transaction may be formed. The transaction defines the specific terms (including the price) and policies by which the item exchanges hands.

The utility an agent gains from any given transaction, is a function of many factors. While for simple products (like CDs), the utility is mostly a matter of price, for more complex products, the purchasing decision generally requires a complex trade-off between a set of preferences. Concentrating on C2C we note that most users in this type of market buy and sell assorted items that are often difficult to describe, and are not easily evaluated. Since in most cases a used item is considered, the value for the buyer will be influenced mostly by the specific functionalities (including attributes like color, size, etc.), quality and the current condition of the product. In addition, the transaction which will eventually be made will include many terms and policies (concerning warranties, return policy, payment policy, delivery time and policy, insurance, etc.). All these terms have values for both sides and can be critical to their buying/ selling decision, regardless of the manner of shopping [15]. Adding reputation and trust factors to the agents' considerations, and keeping in mind that in reference to many features, terms and policies buyers and sellers do not have direct competing interests [12], we assume that the perceived utilities for a buyer and a seller agent from a given transaction are non-correlated. Similar to most other two-sided partnership formation models [8, 40] we assume that the utility of each agent from any given potential transaction can be seen as randomly drawn from a population associated with a specific probability distribution function according to its type. This latter assumption is often justified by the richness of opportunities that may be found in the environment and the complexity of each opportunity. Loyal to the economical principles of market forces as the driving forces that set the terms and conditions of potential partnerships that can be established, we can expect each agent to face a similar distribution of partnership utilities.

The interaction between buyer and seller agents may involve bargaining, however we are interested in the results of the bargaining process (i.e., the terms and policies that define the transaction) rather than the bargaining protocol itself. Therefore bargaining is just another factor contributing to the variance in outcomes and the existence of utilities distribution.

We assume that agents, while ignorant of the utility associated with partnering with specific individual agents of the other type, are acquainted with the overall utility distribution functions[6] and that this distribution remains constant over time. Similarly, the interaction with other agents does not imply any new information about the environment structure.

Before specifying our assumptions concerning the agent's parallel search capability, we would like to re-emphasize that all the assumptions made until this point concerning the utilities distribution are standard and appear in most two-sided search literature cited in this paper (e.g., [8, 40]).

While traditional two-sided partnership formation models assume pure sequential search [8], where each buyer agent is acquainted with only one seller agent in a search stage, in our model we integrate the capability to maintain parallel interactions. Specifically, we assume that agents of the buyer type can consider parallel interactions with several sellers (interested in selling an item similar to the one they seek to buy) at each search round, whereas sellers cannot control the intensity of their interactions. This suits the electronic marketplace domain since in current C2C markets sellers are usually approached by buyers and do not approach buyers in a proactive manner (e.g., buyers and sellers on eBay.com). Still, we do not completely ignore the

[6] There are several methods by which an agent can be acquainted with this distribution function: past experience, bayesian update through sampling, etc.

option in which both sides in the search model use parallel interactions and we briefly discuss some of the aspects of such a model in section 4.2.

We assume that the agents' operation in the eMarketplace, and in particular their search process, is associated with some cost representing the resources they need to consume for their operation. Seller agents interact with one buyer agent at a time, thus their search cost per search stage is fixed. Buyer agents' search cost is more complex and is assumed to combine fixed and non-decreasing variable components, where the variable cost is a function of the number of interactions maintained in parallel. We assume that the marginal cost from interacting with an additional seller agent in parallel is non-decreasing. This is simply because there are some fixed cost components associated with any additional interaction (e.g., communication costs, processing costs [37]).

We assume that the agents' utility from transactions, as well as the resources required for maintaining the search, can be measured on a similar scale. Thus the total search utility can be obtained by subtracting the search "costs" from the perceived utility for any given transaction. Similar to most of the other two-sided partnership formation models [8, 40] we assume that the costs are common to all agents of a given type.[7] This assumption well suits MAS environments since the agents are principally technically similar or supplied by the same market maker (unlike in environments populated with people). A model where different agents of the same type use different search cost structures may also be considered. In this case we obtain a set of sub-types that can be integrated into the appropriate equations given in the analysis section along with their distribution in the general population. For simplification, in this paper we use the analysis where all the agents of a specific type (e.g. buyer agents) share the same cost structure, which is applicable for most markets where agents are supplied to the users by the market maker.

After reviewing and evaluating the perceived utility in a potential partnership each agent makes a decision whether to commit to it or reject it. A transaction (which is the result of the partnership) takes effect only if both agents are willing to commit to it. For simplicity we assume that the agents use synchronous interactions thus if the buyer agent decides to encounter several seller agents in parallel, it commits (if at all) to its "best opportunity", i.e., the one with the highest utility, and rejects the rest of the agents with whom it interacted. A model based on a-synchronous interactions in which the buyer agent considers committing to more than a single opportunity in the same search stage is briefly introduced in section 4.1. If a dual acceptance is not reached, both agents resume their search in a similar manner (with the same cost structures).

Since the agents are self-interested, their goal is to maximize their total search utility (defined above as the perceived utility from the partnership they eventually form minus the accumulated search costs). Therefore, upon meeting a potential partner at any given stage of the search, the agent's problem is to decide whether to form a partnership with this agent or to continue its search. If it commits to a partnership with that agent and that agent commits to the partnership as well (i.e. "dual commitment") then the agent obtains an immediate gain (the expected utility from the partnership). Otherwise, the agent needs to continue searching, bearing additional search costs. In the latter case the agent's future expected utility will be derived from the benefit future interactions might offer, as well as the encountered agents' willingness

[7] One major factor influencing the search cost is the environment opportunities liquidity and volatility, thus the search cost is usually common to most agents in the environment.

to form a partnership. Therefore, an agent's decision of whether to join the current possible partnership depends on the strategies set by the other agents. In a similar manner the other agents' strategies depend on the decision of the agent. Thus, we are seeking strategies that are in equilibrium.

Lastly, we assume that utilities from a partnership are immediate (e.g., payments versus shipment of the product). Furthermore, the nature of eCommerce and agents technology suggest fast search stages and short interaction sessions. Thus even though the agents are not limited by a decision horizon or by the number of search rounds, the probability that an agents' overall search will last more than a few hours/days is negligible. Therefore, in our model we see no need in discounting utilities over time. As we briefly discuss in section 3, the inclusion of a discounting factor does not affect the analysis methodology.

2.2 Assumptions Summary

Before formally introducing the model, we recap the assumptions used and correlate them to the different entities combining our model:

1. **Environment** - populated with numerous self-interested agents;

2. **Agents** - each agent can be of one of two types (e.g., buyers and sellers); benefits from forming a partnership with a single agent of the opposite type; has no a-priori valuation concerning the utility that can be obtained by partnering with specific agents of the other type; can learn the utility of a partnership with a specific agent by interacting with it; is acquainted with the overall utility distribution functions of forming partnerships in the environment; incurs costs when interacting with agents of the opposite type; attempts to maximize its total search utility (defined as the perceived utility from the partnership it eventually forms minus the accumulated search costs);

3. **Partnerships** - form only if both agents commit to them; are associated with non-correlated utilities of the two agents forming them, according to their type;

4. **The Search Process** - each agent interacts randomly with others to learn the utility of partnerships with them; agents of one specific type (buyers) can consider parallel interactions with several agents of the other type (sellers) at each search round, whereas agents of the other type (sellers) cannot control the intensity of their interactions;

5. **The Interaction Protocol** - at the end of the interaction each agent decides whether to accept the partnership with the other agent or reject the partnership; if a dual acceptance is not reached, both agents resume their search; whenever parallel interactions are used, the agent uses a synchronous interaction protocol and commits (if at all) to the partnership associated with the highest utility, while rejecting the rest of the agents with whom it interacted;

6. **Utilities and Costs** - the utility of each agent from any given potential transaction can be seen as randomly drawn from a population associated with a specific probability distribution function according to its type; utilities are immediate and do not need to be discounted; the cost of interacting with other agents is similar to all agents of the same

type; when interacting with several agents in parallel, the marginal cost of interaction with any additional agent is non-decreasing; the cost of search and the utility from a partnership can be measured on the same scale;

2.3 Problem Formulation

We use U^s and U^b to denote the seller and buyer agents' utility perceived from a specific transaction between them (respectively). These utilities can be seen as randomly drawn from a population with p.d.f. $f^s(U^s)$ and c.d.f. $F^s(U^s)$ for the seller agent and $f^b(U^b)$ and $F^b(U^b)$ for the buyer agent, $(0 \leq U^s < \infty, 0 \leq U^b < \infty)$. The buyer agents' search cost of sampling N seller agents is denoted $C(N)$ (satisfying $\frac{dC(N)}{dN} > 0$ and $\frac{d^2C(N)}{dN^2} \geq 0$). The seller agents' search cost is fixed per any search round, denoted by c. We use U^b_{best} to denote the utility associated with the "best opportunity" among the potential opportunities available to an agent of the buyer type upon interacting with N agents of the seller type in a given search round, i.e., $U^b_{best} = max(U^b_{(1)}, \ldots, U^b_{(N)})$.

The problem of each seller agent is to find a strategy $S_{seller} : U^s \rightarrow \{accept, reject\}$ that will map the perceived utility of each potential interaction to a decision of whether to commit to the proposed transaction or reject it. As for the buyer agent, its strategy is a little more complex. The buyer agent needs to set a strategy, given the value of the best opportunity found in the current search round, that defines both whether to commit to an opportunity or reject it and with how many seller agents it wants to interact in parallel over the next search round (if at all). Thus the buyer agent's strategy can be specified as $S_{buyer} : U^b_{best} \rightarrow \{accept, N\}$ where $accept$ suggests committing to the best opportunity in the current search stage (yielding a utility U^b_{best}) and N is the number of seller agents to interact with next if the current set of opportunities is rejected.

Before concluding the problem definition section, we formally summarize the two-sided search mechanism detailed above. We divide the description according to the agent's type. Each buyer type agent will be performing the following steps:

loop
  Set a value N
  Locate randomly N seller type agents and initiate interaction with all of them in parallel
  Evaluate the utility $U^b_{(i)}$ from a partnership (transaction) with each agent i, $(i \leq N)$.
  Reject all partnerships with a utility smaller than U^b_{best}
  if Satisfied with utility U^b_{best} **then**
    Commit to a partnership with the agent associated with this utility
    if Other agent commits to the partnership as well **then**
      Terminate the process (terminate search)
    end if
  else
    Reject the agent associated with this utility
  end if
end loop

Each seller type agent will be performing the following steps:

```
loop
  Wait to be approached by a buyer type agent
  Interact with the agent and evaluate the utility U^s from a partnership with it.
  if Satisfied with utility U^s then
    Commit to the partnership
    if Other agent commits to the partnership as well then
      Terminate the process (terminate search)
    end if
  else
    Reject the agent associated with this utility
  end if
end loop
```

For convenience, we have added a table at the end of the paper summarizing all the notations used and their meanings.

3 Analysis

Since the agents are not limited by a decision horizon, and given the fact that the interaction with other agents does not imply any new information about the market structure and that the number of search rounds is not limited, the agents' search strategy is stationary (i.e. an agent will not accept an opportunity it has rejected beforehand, and the number of seller agents with which each buyer agent interacts during a search round, N, will remain constant over time). Similarly, the number of search rounds an agent has already participated in does not affect its optimal strategy, since the cost incurred in earlier search rounds is considered "sunk cost" and do not affect future expected utility from resuming the search, calculated from any decision point and on. This implies a reservation-value based strategy both for the buyer and seller agents.[8] We denote the seller agent's reservation value by x^s and the buyer agent's reservation value by x^b. The agents' reservation-value strategy suggests that a commitment for a potential transaction will be received from the seller agent only if $U^s \geq x^s$ (i.e., if its immediate utility U^s is greater than or equal to its reservation value x^s), and from the buyer agent only if $U^b \geq max(U^b_{(1)}, \ldots, U^b_{(N)}, x^b)$ (i.e., only if the current transaction yields a utility U^b that is the highest utility among the different interactions maintained by the buyer agent at this stage and if it is greater or equal to the reservation value it uses). We use x^b_N to denote the reservation value that maximizes the buyer agent's utility as a function of the number of parallel interactions it maintains during a search round, N, and the strategy used by the seller agents. Similarly, we use x^s_N to denote the reservation value that maximizes the seller agent's expected utility as a function of the reservation value and the number of parallel interactions, N, used by buyer agents.

Our goal is to find a set of equilibrium strategies, i.e., a reservation value x^s_N for the seller agents and a pair (N, x^b_N) for the buyer agents, from which no single agent of any of the types

[8] Notice the reservation value of the search strategy is different from the reservation price usually associated with a buyer or a seller that are not involved in a search. While the reservation price denotes an agent's true evaluation of a specific potential opportunity, the reservation value of a search strategy is mainly a lower bound for the utility from an accepted opportunity, derived from the expected utility optimization considerations.

will have an incentive to deviate, as long as the other buyer and seller agents follow this strategy. For this purpose, we first analyze the strategies that will be used by each agent type, for any given number of parallel interactions, N. This is achieved by understanding how an agent's strategy is affected by changes in the strategies used by agents of the other type. The analysis is followed by an efficient approximation means for extracting the agents' reservation value in those cases where the strategies can not be immediately obtained using direct calculation. Consequently, based on the analysis presented, we manage to suggest an efficient means for finding the number of parallel interactions, N, that is used by the buyer agents in equilibrium and the appropriate equilibrium reservation values both for the buyer and seller agents.

3.1 Agents' Strategies

We begin by formulating the appropriate equations describing each agent's expected utility as a function of the strategy it uses and the strategy used by the other agents' type. Consider a buyer agent using a reservation value x^b, interacting in parallel with N seller agents at each stage of its search. This buyer agent's expected utility, denoted $V^b(x^b)$, when seller agents use a reservation value x^s is given by:

$$V^b(x^b) = E\Big[U^b_{best} \bullet 1[(U^b_{best} \geq x^b) \cap (U^s \geq x^s)] + V^b(x^b) \bullet 1[(U^b_{best} < x^b) \cup (U^s < x^s)] - C(N)\Big] \quad (1)$$

where the term $\bullet 1[(U^b_{best} \geq x^b) \cap (U^s \geq x^s)]$ represents the indicator of the event where the specific buyer agent and its "best" encountered seller agent (in the current search round) have found the perceived utility from a transaction between the two of them to be greater than or equal to their reservation values, resulting in a dual commitment. The term $\bullet 1[(U^b_{best} < x^b) \cup (U^s < x^s)]$ represents the indicator of the complementary event (i.e., where the seller and/or the buyer reject the potential partnership between them). In the latter case, the buyer agent resumes its search in the same manner, i.e., conducts an additional search round in which it interacts with N seller agents using the same reservation-based acceptance rule, hence resulting in an expected future utility which equals $V^b(x^b)$ due to the stationary nature of the problem. Notice that if gains are to be discounted, then the only required change in Equation 1 above is the discounting of the expected future utility $V^b(x^b)$ on the right hand side of the equation.

A similar equation can be formulated to represent the seller agents' expected utility, denoted $V^s(x^s)$, when using a reservation value x^s and given the number of parallel interactions, N, and the reservation value x^b used by the buyer agents:

$$V^s(x^s) = E\Big[U^s \bullet 1[(U^b = U^b_{best}) \cap (U^b \geq x^b) \cap (U^s \geq x^s)] + \\ + V^s(x^s) \bullet 1[(\exists U^b_{(j)} > U^b, j \leq N) \cup (U^b < x^b) \cup (U^s < x^s)] - c\Big] \quad (2)$$

In this case, we need to ensure that U^b (the utility for the buyer from a partnership with the specific seller) is the highest among the utilities from partnering with any of the seller agents with whom it interacts. This is in addition to the requirement that both agents' (the seller's and the buyer's) perceived utilities are equal to or greater than the reservation values they set.

For simplification, we use the notations $F^b_N(x)$, $f^b_N(x)$ and $E[U^b_N]$ to denote the c.d.f., p.d.f. and the mean of the maximum utility for the buyer agent when interacting with N seller agents in parallel, respectively. Notice that these values can be calculated using the standard utility c.d.f and p.d.f functions ($F^b(x)$ and $f^b(x)$) as follows:[9]

$$F^b_N(x) = (F^b(x))^N \qquad f^b_N(x)\frac{dF^b_N(x)}{dx} = NF^b(x)^{N-1}f^b(x) \qquad E\Big[U^b_N\Big]=\int_{y=0}^{\infty}(1 - F^b_N(y))dy \quad (3)$$

Using the above notations, we attain (see Appendix A for more details):

$$V^b(x^b) = \frac{(1 - F^s(x^s)) \int_{y=x^b}^{\infty} y f^b_N(y)dy - C(N)}{(1 - F^b_N(x^b))(1 - F^s(x^s))} \tag{4}$$

In order to simplify the flow of the paper, we will present only the equations associated with the buyer type agents. Unless stated otherwise, similar modifications for the seller type agents are given in Appendix B or can be extracted using methods similar to those we use for the buyer type agents' equations.

Equation 4 and its sellers' modification can be used by each agent to calculate its expected utility, when using different reservation value strategies, given the search cost parameters and the strategy used by the agents of the opposite type. From this equation we can derive an agent's reaction to changes in the other agents' strategies, towards a complete equilibrium analysis. Notice that Equation 4, as well as the rest of the following analysis, is also applicable for the traditional pure sequential two-sided search [8], simply by using $N = 1$.

Continuing our analysis, we point to an immediate result from Equation 4:

$$\lim_{x^b \to \infty} V^b(x^b) = -\infty \quad ; \quad \lim_{x^b \to 0} V^b(x^b) = E[U^b_N] - \frac{C(N)}{1 - F^s(x^s)} \tag{5}$$

The content of Equation 5 is intuitive: if the reservation value x^b is very large, the chances of obtaining a utility greater than this reservation value from a given search round are small. Thus, repeated search rounds must be taken, leading to an overall low utility due to the accumulated search costs. If, on the other hand, the reservation value, x^b, is very small, almost surely a potential transaction suggesting a sufficient utility can be obtained during the first search round.

Theorems 1-2 below, suggest several additional important properties of the agents' expected utility function, that are used later for designing the algorithms for extracting the agents' strategies. Appropriate modifications of these theories for the seller agents are given in Appendix B.

Theorem 1 *The expected utility function $V^b(x^b)$ is quasi concave, with a unique maximum obtained at point x^b_N, satisfying:*

$$V^b(x^b_N)=x^b_N \tag{6}$$

Proof: Deriving Equation 4 we obtain:

$$\frac{dV^b(x^b)}{dx^b}=\frac{f^b_N(x^b)(V^b(x^b) - x^b)}{(1 - F^b_N(x^b))} \equiv r(x^b)(V^b(x^b) - x^b) \tag{7}$$

[9] The third equation can be obtained using integration by parts over the expression: $\int_{y=0}^{\infty} y f^b_N(y)dy$.

The value x_N^b for which the above expression equals zero must satisfy $V^b(x_N^b) = x_N^b$. Notice that $f_N^b(x_N^b) > 0$ implies $r(x_N^b) > 0$, hence for x_N^b satisfying $V^b(x_N^b) = x_N^b$:

$$\frac{d^2V^b(x_N^b)}{dx_N^{b\,2}} = r^{'}(x_N^b)(V^b(x_N^b) - x_N^b) + r(x_N^b)(V^{b^{'}}(x_N^b) - 1) < 0 \quad (8)$$

Thus $V^b(x_N^b)$ (and in the same manner $V^s(x_N^s)$) is quasi concave with a unique maximum.□

Equality 6 is very common in models in which agents use reservation-value based strategies [26]. It suggests that the expected utility when using the optimal reservation value equals the optimal reservation value. Intuitively, this can be explained by the fact that the agent's optimal reservation value is the point where it becomes indifferent to the selection between the utility that can be obtained from a transaction and the utility associated with continuing the search. Similarly, if we were about to use a discounting factor λ for gains, Equation 6 would transform into: $V^b(x_N^b) = \lambda x_N^b$.[10]

In order to illustrate the dynamics of changes in the agent's policy we use the following synthetic artificial environment.[11]

Environment 1 *The environment contains numerous buyer and seller agents, where each interaction between any buyer and any seller agents produces utilities drawn from a triangular distribution function, i.e.:*[12]

$$f(x) = \frac{2}{U_{upper}} - \frac{2x}{(U_{upper})^2} \qquad F(x) = \frac{2x}{U_{upper}} - \frac{x^2}{(U_{upper})^2} \qquad (0 \leq x \leq U_{upper})$$

where in our case $U_{upper} = 100$ and the search costs are $C(N) = 2 + 0.5N$ and $c = 2.5$ for the buyer and the seller agents, respectively.[13] *The environment can be used both for experimenting with the traditional sequential (single interaction at a time) search and the parallel-based search.*

Figure 1, illustrates the agents' expected utility as a function of their reservation value in two settings (here it is assumed that the agents of the other type use the value 30 as their reservation value). In the first setting, all the agents use the pure sequential search ($N = 1$). The middle curve depicts the expected utility of any of the agents in this scenario as a function of the reservation value used (the horizontal axis). Since both the buyer and the seller agents use a sequential search, the curve describes a utility function that is common to all the agents. In the second setting, buyer type agents use the new parallel search method, sampling $N = 4$ seller agents in each search round. In this scenario buyer type agents have the incentive to use the new technique (represented by the upper curve) since their utility increases for any reservation value they use in comparison to the pure sequential search. Similarly, the expected utility for

[10] From this point and on the introduction of the appropriate utility-discounting modification to each equation used is straightforward and fits into the analysis methodology presented.

[11] As we show in the following paragraphs the transition of buyer agents to the parallel interaction technique is inevitable in many environments and non-questionable given the benefits encapsulated in this mechanism. Therefore, our goal with this example is mainly to demonstrate the different theorems and analysis that we introduce throughout the paper. The selected synthetic artificial environment is rich enough for this purpose.

[12] This distribution function can be associated with most electronic marketplaces. It reflects a high probability to draw an opportunity producing a low utility, and vice versa.

[13] Thus when buyer agents use N=1, all agents' search cost structures are symmetric.

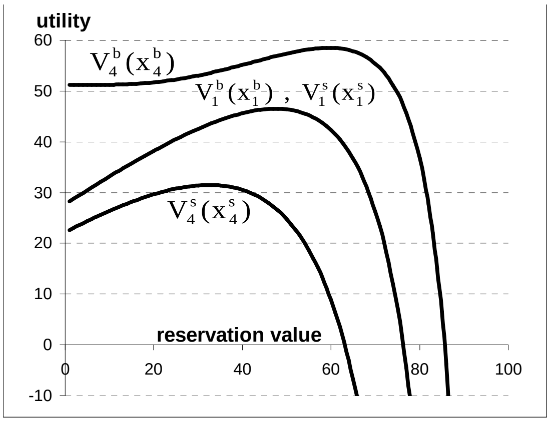

Figure 1: Agent's expected utility as a function of its reservation value in Environment 1

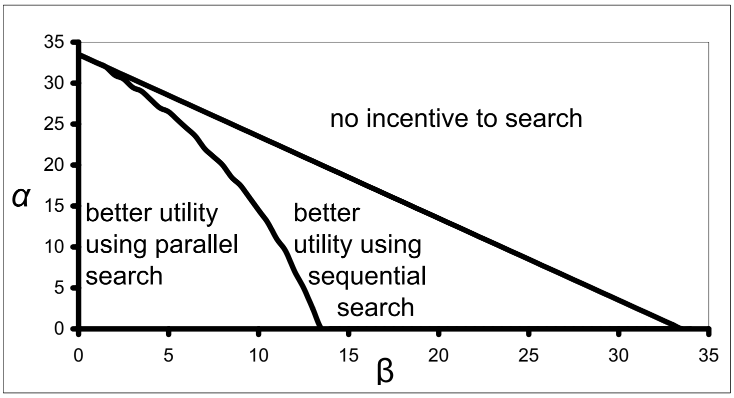

Figure 2: Agents' incentive to use parallel search (in comparison to sequential search).

the seller type agents (represented by the lower curve) always decreases when the buyer agents adopt the new method.

The incentive for the buyer agents to use the combined parallel search technique is strong. Any single buyer agent will prefer to use more than a single interaction during a search round, if it finds the expected utility to be higher in this manner. Figure 2 demonstrates this phenomena using Environment 1. The cost functions used for this example are: $C(N) = \alpha + \beta N$ and $c = \alpha + \beta$. A two-sided search in this environment will take place only if the expected utility for the agents when using the pure sequential equilibrium strategy is positive. The set of α and β combinations that guarantee a positive utility (when using the equilibrium set of strategies with $N = 1$) is represented by the bottom triangular area in the graph. Out of this area, we have isolated (on the bottom left side) all combinations of α and β values where a buyer agent can increase its expected utility by deviating from such a pure sequential strategy (assuming all other agents' strategies are sequential) to a parallel search strategy (i.e., using $N > 1$). We learn from the graph that buyer agents have an incentive to deviate from the traditional pure sequential search strategy for many plausible combinations of α and β values. Furthermore, the advantage of the new technique is mostly revealed in combinations of small α and β values (in comparison to the average utility from a partnership), which characterize most MAS applications. A simple tool (applicable for any distribution function) for checking if buyer agents have an incentive to deviate from the pure sequential search technique to the parallel one is given in Theorem 5 (in section 3.5).

Next, we introduce a theorem that enables us to extract the agents' optimal reservation value as a function of the reservation value used by agents of the other type. Notice that our aim, at the current stage, is to find the reservation value that maximizes the agent's utility given the strategy used by other agents' type, rather than finding the equilibrium strategies.

Theorem 2 *Given the reservation value that was set by the seller agents, x_N^s, the buyer agents' optimal reservation value when using N parallel interactions, x_N^b, satisfies:*

$$C(N) = (1 - F^s(x_N^s))(E[U_N^b] - \int_{y=0}^{x_N^b} (1 - F_N^b(y))dy) \tag{9}$$

Proof: Deriving the expected utility given in Equation 4, setting it to zero, and using integration by parts for calculating $\int_{y=x_N^b}^{\infty} y f_N^b(y)dy$, we finally obtain Equation 9.□

From Theorem 2 we conclude that the buyer agents' optimal reservation value (and thus the expected utility for these agents, based on Theorem 1) decreases as $C(N)$ increases. This also has an intuitive explanation: when the search costs increase, the agent becomes less selective, reducing its reservation value. Secondly, we can conclude from Theorem 2 that the buyer agents' optimal reservation value (and thus the total utility for the agent), given the seller agent's reservation value, decreases as x_N^s increases. Similar results can be obtained for the seller (see Appendix B).

3.2 Approximation Technique

Both Equations 6 and 9, and their appropriate modifications for the seller agent, can be used for calculating the optimal reservation values of any agent type in the search, given the reservation values used by the agents of the opposite type. However, for some distribution functions (e.g., normal distribution function) it is impossible to extract x_N^s and x_N^b using direct immediate calculations (see the example given in Appendix B).[14] Fortunately, the characteristics of the optimal strategies (given the other agents' reservation values) as proven in theorems 1-2, enable us to use binary search as an efficient means for approximating these values up to any required precision level. In order to apply binary search here we first need to find an interval that bounds the optimal reservation value of the agent and a condition by which we can tell whether any of the specific values in this interval are greater or smaller than the optimal reservation value. Notice that in most scenarios the distribution functions of U^s and U^b are finite (assuming the utility of the person represented by an agent from a specific exchange is finite) and can be used as a bounding interval for the optimal reservation value of the agents. Still, in the following Proposition 1 we supply a bounding interval and necessary conditions for binary search within

[14]The case of non-integrated terms in Equation 9 (due to the complexity of a specific distribution function) can be resolved by using specific function-dependent approximation techniques. For example, solving Taylor series expansion, or using the Trapezoidal Rule and Simpson's Rule. In this section we introduce a general approximation algorithm (i.e., distribution-independent) based on binary search.

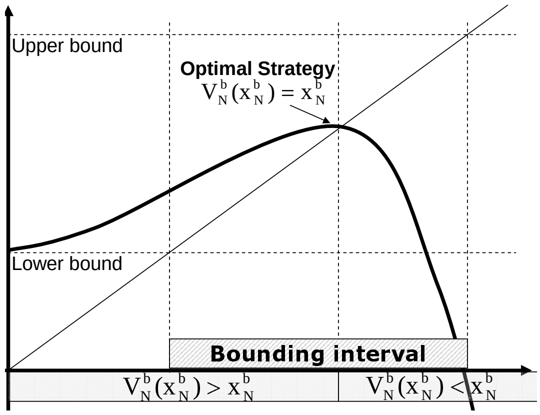

Figure 3: Specific characteristics of an agent's expected utility function $V^s(x^s_N)$

this interval that are applicable even for distribution functions that are not necessarily defined over a finite interval.

Proposition 1 *(a) The values* x^b_{N+} *satisfying* $(1 - F^s(x^s_N)) \int_{y=x^b_{N+}}^{\infty} y f^b_N(y)dy = C(N)$ *and* $x^b_{N-} = 0$*, can be used as upper and lower bounds, respectively, for the buyer agents' optimal reservation value (given the seller agents' reservation value). Similar bounds can be found for the sellers' optimal reservation value; (b) Within the interval defined above, any value* x^b *for which* $V^b(x^b) > x^b$ *is smaller than the optimal reservation value and vice versa.*

Proof: For the proof we make use of Figure 3 which sketches the general shape of any of the agents' expected utility as a function of the reservation value used.[15] Taking the buyer type agents, as an example, the expected utility for the case where $x^b = 0$ is $E[U^b_N] - \frac{C(N)}{1-F^s(x^s)}$ (see Equation 5). From this point and on the expected utility increases as x^b increases, reaching a global maximum at the point x^b_N where $V^b_N(x^b_N) = x^b_N$ and decreases beyond that point (derived from Theorem 1 that indicates the concavity of the function and its limit given in Equation 5 for $x^b \to \infty$). Since $V^b_N(0) > 0$ (see Appendix C for the analysis of the case where $V^b(0) < 0$), the expected utility $V^b(x^b)$ calculated using Equation 4 is always greater than x^b for $x^b < x^b_N$ and always smaller than x^b for $x^b > x^b_N$. Now all we need to prove is that x^b_N is in the bounding interval specified in the theorem. For the upper bound, we note that a reservation value x^b_{N+} that satisfies the condition in part (a) of the theorem always yields $V^b(x^b_{N+}) = 0 < x^b_{N+}$ when used in Equation 4. Therefore, given the antecedent part of the proof, x^b_{N+} is an upper bound for the optimal reservation value, x^b_N. The lower bound, $x^b_{N-} = 0$, is obvious. □

Therefore, an algorithm for extracting the optimal reservation value, x^b_N, for the buyer agents (or x^s_N for the seller agents) should first check the value of $V(0)$ using Equation 4. If this latter value is negative then $x^b_N = 0$ (see the analysis in Appendix C). Otherwise, the algorithm should set the appropriate bounding values (according to Proposition 1) and conduct a binary search in this interval, each step checking if the current value x^b is greater or smaller than $V^b(x^b)$. The interval should be trimmed to include all values smaller than x if $V^b(x^b) > x^b$ and

[15] A unique scenario, in which the expected utility function strictly decreases and never satisfies $V^b(x^b_N) = x^b_N$, is analyzed in Appendix C.

vice versa. The algorithm should stop when reaching a value x^b for which $|V_N^b(x^b) - x^b| \leq \rho$, where ρ is the required precision level. These steps are summarized in Algorithm 1.

Algorithm 1 An algorithm for calculating (using approximation) the optimal reservation value x_N^b for a buyer agent, given the strategies used by the other agents and the number of parallel interactions it uses

Input: ρ - required precision level; N - number of parallel interactions used by the buyer agent; x_N^s - the reservation value used by the seller type agents.
Output: x_N^b - precise approximation for the buyer agent's optimal reservation value.

if $V_{equation(4)}(0) < 0$ **then**
  return 0
end if
Calculate the values of x_{upper} and x_{lower} according to Proposition 1
while $| V_{equation(4)}(x) - x |> \rho$ **do**
  Set $x = (x_{lower} + x_{upper})/2$
  if $V_{equation(4)}(x) > x$ **then**
    Set $x_{upper} = x$
  else
    Set $x_{lower} = x$
  end if
end while
return x

The notation $equation(i)$ is used in the algorithm to denote the calculation of the parameter using Equation i.

Proposition 2 *Given the number of parallel interactions buyer agents use, N, a reservation value of the seller, x_N^b and a required precision level, ρ, Algorithm 1 returns a reservation value for a buyer agent that is close to the optimum within ρ. The complexity of the algorithm is $O(ln\frac{x_{N+}^b}{\rho})$.*

Proof: According to Proposition 2, the algorithm always handles the interval that bounds the optimal reservation value. Since the algorithm is based on binary search, its complexity is given by $O(ln\frac{x_{N+}^b}{\rho})$. □

3.3 Finding a Stable Set of Reservation Values

The equilibrium in our model can be described by a set (N, x_N^b, x_N^s) where a single buyer agent cannot gain a better utility by changing N and/or x_N^b which it uses and a seller agent cannot gain a better utility by changing x_N^s (as long as the other agents do not change their strategies). Using the analysis given in the previous section, we can now combine the reactions of both types of agents to changes in the other agents type reservation value, towards equilibrium.

Recall that an important result from Equation 9, is that an agent's optimal reservation value decreases as the reservation value used by the agents of the opposite type increases, i.e., for any specific N we obtain:

$$\frac{dx_N^b}{dx_N^s}, \frac{dx_N^s}{dx_N^b} < 0 \qquad \lim_{x_N^s \to \infty} x_N^b = \lim_{x_N^b \to \infty} x_N^s = 0 \tag{10}$$

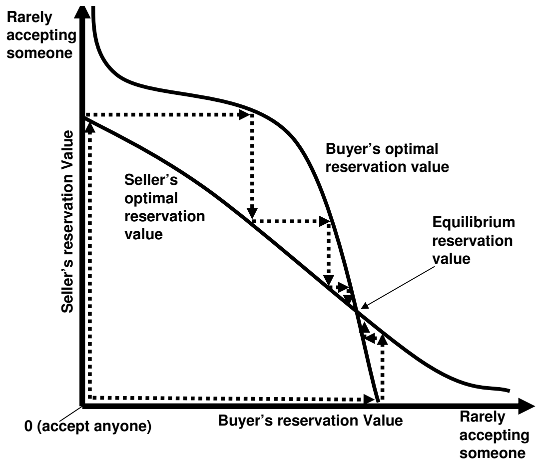

Figure 4: Agents' reaction curves - the buyer agent's optimal reservation value as a function of the reservation value used by the seller agents (vertical axis) and the seller agent's optimal reservation value as a function of the reservation value used by the buyer agents (horizontal axis).

This behavior is illustrated in Figure 4, which reflects characteristics similar to those described in [8] for the pure sequential two-sided search. Nevertheless, while in [8] any point of interception between the two curves is an equilibrium point, in our model every such point is merely a potential "suspected" equilibrium point. If the buyer agents can only change their reservation value, then each such point is an equilibrium point for the specific N value used by the buyer agents. Nevertheless, since buyer agents can control the number of seller agents with whom they interact in parallel then these intersection points can merely be considered "suspected" equilibrium points and the determination of whether these points are in equilibrium requires further validation as described in the following paragraph.

From Equation 10 we conclude that at least one "suspected" equilibrium exists for each N value (in the extreme case, we obtain a set of strategies where agents of a specific type or of both types accept any agent of the opposite type). In some environments we have a single "suspected" intersection point (for example, for the uniform distribution function). However, theoretically, a general distribution function might produce several equilibria points with uncertainty regarding the identity of the one that will eventually be used. None of these "suspected" equilibria dominates the other for both agent types (buyer agents and seller agents), i.e., none of the "suspected" equilibria Pareto dominates any of the others. This can be observed by recalling the fact that $V^b(x^b_N) = x^b_N$ for both agents (Equation 6) and given the structure of the two curves in Figure 4. An additional important result is that the total number of "suspected" equilibrium strategies sets for any given N is odd. This, again, is derived from the unique structure of the two curves described in Figure 4.

In order to find all these "suspected" equilibrium sets of strategies, for any given N value, we can use the following Algorithm 2.

The algorithm operates in a recursive form (see steps 17-18). It starts with an initial set of "suspected" equilibrium sets of strategies, $Q = \{(x^b_N, x^s_N)_1, \ldots, (x^b_N, x^s_N)_k\}$, that were found in earlier recursive executions of this algorithm for the specific N value that is used. It adds to

Algorithm 2 $FindSuspect(Q, \rho, N, (x_{upper}, x_{lower}))$ - An algorithm for finding "suspected" equilibria when the buyer agents use N parallel interactions at each search round.

Input: Q - a set of "suspected" equilibria strategies; ρ - required precision level; N - number of parallel interactions used by the buyer agent; (x_{upper}, x_{lower}) - an interval in which the buyer agents' reservation values of new "suspected" equilibria should reside.

Output: Q - an extended set of "suspected" equilibria strategies.

1: Set $x^b = x_{upper}$
2: Set $x^s_{store} = x^b_{store} = 0$
3: **for** ($i = 1$ to 2) **do**
4: $diff = \rho + 1$
5: **while** $diff > \rho$ **do**
6: Set $x^s = x^s_{equation(37)}(x^b)$
7: Set $x^b = x^b_{equation(9)}(x^s)$
8: Set $diff = \mid x^s - x^s_{store} \mid + \mid x^b - x^b_{store} \mid$
9: Set $x^s_{store} = x^s$; $x^b_{store} = x^b$
10: **end while**
11: **if** $(x^b, x^s) \notin Q$ **then**
12: add (x^b, x^s) to Q
13: **end if**
14: Set $x^b = x_{lower}$
15: **end for**
16: **if** found different sets (x^b, x^s) in both executions of steps 4-13 **then**
17: $Q = FindSuspect(Q, \rho, N, (x_{lower}, \frac{x_{upper}+x_{lower}}{2})$
18: $Q = FindSuspect(Q, \rho, N, (\frac{x_{upper}+x_{lower}}{2}, x_{upper})$
19: **end if**
20: **return** Q

this set new "suspected" equilibrium sets upon finding them. Thus for its initial execution we can use $Q = \emptyset$ as an input. On each of its recursive executions it operates on a specific interval of buyer agents' reservation values. Therefore, for its initial execution we can use the interval $(0, x^b_N)$ where x^b_N is the optimal reservation value of the buyer type agents when seller agents use a reservation value zero (x^b_N can be found using either a direct calculation through Equation 9 or Algorithm 1 that was given in the previous subsection). Since x^b_N reaches its maximum value when $x^s = 0$ (see Equation 10) then there will be no set of strategies in equilibrium in which buyer agents use a reservation value greater than this latter x^b_N value. Therefore all suspected equilibrium sets are within the interval with which the algorithm begins. In its main loop (steps 5-10) the algorithm finds a set of strategies "suspected" to be in equilibrium, by sequentially changing the reservation values used by the different agent types, where at each stage one type sets its optimal reservation value given the last reservation value set for the other type. This calculation method always converges to a point of intersection of the two curves as depicted in Figure 4 (i.e., "suspected" equilibrium set) since each agent's reservation value is sequentially increased/decreased (as a reaction to the changes in the other agent's reservation value) in a decreasing rate, in each subsequent stage of the process. If only a single equilibrium set exists in the specified interval, then when executing steps 5-10 from both ends of the interval the same set of reservation values is reached, thus there is no point in searching for further "suspected" strategy sets in this interval. Otherwise, if a new "suspected" set of strategies is found, then we divide the interval into two parts and recursively activate the algorithm for any of these two new intervals (steps 17-18). This is due to the fact that we do not know the

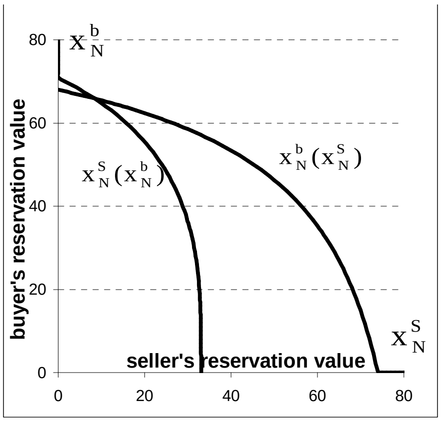

Figure 5: Finding "suspected" equilibrium sets of strategies (N=4)

number of "suspected" equilibrium sets that might exist in this interval. In any case, any new "suspected" equilibrium sets of strategies are added to the set Q (steps 11-13).

Proposition 3 *Algorithm 2 will reach all suspected equilibria strategy sets. The number of recursive executions of itself that will be initiated is bounded by $Z * ln(\frac{d}{\rho})$, where d is the interval given in its most initial execution and Z is the number of suspected equilibria in the specific environment.*[16]

Proof: Further recursive executions of the algorithm occur only if the interval contains "suspected" equilibrium sets (steps 17-18), in which case the interval is divided into two equal sub-intervals. If the given interval does not contain "suspected" equilibrium sets (i.e., both executions of steps 4-13 end up with the same set of strategies) then no further executions take place for this interval. Therefore, the number of recursive executions is equivalent to a binary search (in an interval smaller than d) for any "suspected" set of strategies that exist (i.e., for any intersection of the two curves). □

Figure 5 illustrates the changes in the strategies of the two-types of agents when using Environment 1 (where buyer agents use $N = 4$). From this figure we can clearly see that the singular intersection point is where buyer type agents use $x^b_4 = 65.6$ and seller type agents use $x^s_4 = 9.2$. This is in comparison to the (46.5,46.5) equilibrium reservation values, associated with the traditional sequential two sided search model (for $N = 1$), obtained from the middle curve in Figure 1 by calculating $V^b(x^b_1) = x^b_1$. Recall that in Theorem 1 we obtained $V^b(x^b_N) = x^b_N$ (and similarly for the seller type agents). Thus in the former scenario, the buyer type agents' revenue increases at the expense of the seller type agents. This is always true since the buyer agents become more selective, thus the probability for the seller agents to be accepted in a given encounter decreases (and the increase in the number of search rounds the seller agents need to execute results in increased aggregated search costs).

[16]The actual number of times the code is executed is smaller since the search is performed in intervals that decrease exponentially on each execution.

3.4 Finding Equilibria Strategies

Now that agents of both types are capable of calculating their strategies, given the number of parallel interactions set by the buyer agents, N, we proceed to find the equilibria strategies. As noted in the previous subsection, the agents can identify suspected equilibrium strategies set for each N value. However in order to guarantee that this is a set of equilibria strategies, we need to make sure no single buyer agent has an incentive to deviate from its specified strategy by using a different strategy $(N', x_N^{b'})$ while all other buyer agents use strategy (N, x_N^b) and all other seller agents use strategy x_N^s. Fortunately, we do not need to check the stability of the set (N, x_N^b, x_N^s) for any $N' \neq N$ value. The following Theorem 3 suggests that we only need to check if a single buyer agent has an incentive to deviate towards using strategies $(N+1, x_{N+1}^b)$ and $(N-1, x_{N-1}^b)$ in order to declare a set of strategies as equilibrium.

Theorem 3 *(a) If* $V_N^b(x_N^b) > V_{N+1}^b(x_{N+1}^b)$ *then* $V_N^b(x_N^b) > V_{N+k}^b(x_{N+k}^b)$, $\forall k > 1$*; (b) If* $V_N^b(x_N^b) > V_{N-1}^b(x_{N-1}^b)$ *then* $V_N^b(x_N^b) > V_{N-k}^b(x_{N-k}^b)$, $\forall 1 < k < N$.

Proof: We first prove part (a) of the theorem. Assume otherwise, i.e., assume $V_N^b(x_N^b) > V_{N+1}^b(x_{N+1}^b)$ and $V_{N+1}^b(x_{N+1}^b) < V_{N+2}^b(x_{N+2}^b)$. Therefore, given Equation 9, we obtain:

$$C(N) + C(N+2) - 2C(N+1) = \tag{11}$$
$$= (1 - F^s(x_N^s))\Big(\int_{y=x_N^b}^{\infty}(1 - F_N^b(y))dy + \int_{y=x_{N+2}^b}^{\infty}(1 - F_{N+2}^b(y))dy - 2\int_{y=x_{N+1}^b}^{\infty}(1 - F_{N+1}^b(y))dy\Big)$$

Since $C(N)$ is an increasing linear (or convex) function of N then $C(N)+C(N+2)-2C(N+1) \geq 0$. Also, from the assumption used in the prelude of the proof, we know (using Theorem 1) that: $x_N^b > x_{N+1}^b$ and $x_{N+1}^b < x_{N+2}^b$, thus:

$$0 < \int_{y=x_{N+1}^b}^{\infty}(1-F_N^b(y))dy + \int_{y=x_{N+1}^b}^{\infty}(1-F_{N+2}^b(y))dy - 2\int_{y=x_{N+1}^b}^{\infty}(1-F_{N+1}^b(y))dy = -\int_{y=x_{N+1}^b}^{\infty} F^b(y)^N \big(1-F^b(y)\big)^2 \tag{12}$$

where the last equality is derived using the substitution $F_i^b(y) = (F^b(y))^i$ as given in Equation 3. The result of Equation 12 is a contradiction, thus the assumption made at the beginning of the proof is invalid. Thus $V_N^b(x_N^b) > V_{N+1}^b(x_{N+1}^b)$ necessarily results in $V_{N+1}^b(x_{N+1}^b) > V_{N+2}^b(x_{N+2}^b)$. Therefore, if $V_N^b(x_N^b) > V_{N+1}^b(x_{N+1}^b)$ then the value of $V_k^b(x_k^b)$ ($\forall k > N$) decreases as k increases, thus $V_N^b(x_N^b) > V_{N+k}^b(x_{N+k}^b)$, $\forall k > 1$. A similar proof can be written for part (b) of the theorem. □

Therefore, in order to find an equilibrium in a given environment, one needs to find the "suspected" equilibrium strategy sets for each N value (as explained in subsection 3.3), and then check if stability holds for these sets by evaluating the buyer type agents' utility when shifting to $N+1$ or $N-1$ parallel interaction strategies. Nevertheless, this process requires an upper bound for N (for a lower bound we can use $N = 1$). The following theorem suggests such an upper bound.

Theorem 4 *The number of parallel interactions, N^*, that is used by buyer agents in equilibrium is bounded (from above) by any N that satisfies:*

$$N > max\big(\frac{E[U^s]}{c}, \frac{E[U^b_N] + C(2) - 2C(1)}{C(2) - C(1)}\big) \tag{13}$$

Proof: First, we will prove Lemmas 1-3. The first lemma suggests that the expected maximum utility of a sample increases as the sample size increases (however in a decreasing rate). This lemma is then used in the proof of Lemmas 2-3, that specify the condition under which both seller and buyer agents will accept any agent they come across.

Lemma 1 *(1)* $\frac{dE[U^b_N]}{dN} > 0$ *(2)* $\frac{dE[U^b_N]}{d^2N} < 0$

Proof:
(1) Formulating the explicit expression $E[U^b_{N+1}] - E[U^b_N]$, we obtain:

$$E[U^b_{N+1}] - E[U^b_N] = \int_{y=0}^{\infty} (F^b(y)^N - F^b(y)^{N+1})dy \tag{14}$$

which is always positive.
(2) Formulating the explicit expression $E[U^b_{N+2}] - 2E[U^b_{N+1}] + E[U^b_N]$, we obtain:

$$E[U^b_{N+2}] - 2E[U^b_{N+1}] + E[U^b_N] = -\int_{y=0}^{\infty} F^b(y)^N(1 - F^b(y))^2 dy \tag{15}$$

which is always negative. □

Lemma 2 *Seller agents will commit to any partnership if the buyer agents use $N > \frac{E[U^s]}{c}$ parallel interactions.*

Proof: Substituting $N > \frac{E[U^s]}{c}$ in the appropriate seller's modification for Equation 4 (see Equation 33 in Appendix B) we obtain:

$$V^s(x^s) < \frac{(1 - F^b(x^b)^N) \int_{y=x^s}^{\infty} yf^s(y)dy - E[U^s]}{(1 - F^b(x^b)^N)(1 - F^s(x^s))} = \tag{16}$$
$$= \frac{-\int_{y=x^s}^{\infty} yf^s(y)F^b(x^b)^N dy - \int_{y=0}^{x^s} yf^s(y)}{(1 - F^b(x^b)^N)(1 - F^s(x^s))}$$

Now notice that the right hand side of Equation 16 above is negative for any x^s value. Therefore $V^s(0) < 0$, and according to Theorem 8 (given in Appendix C), $V^s(x^s)$ is a strictly decreasing function of x^s. Thus the seller agents' strategy in this scenario would be to use $x^s = 0$, i.e., accept any partnership. □

Lemma 3 *Buyer agents will use a reservation value $x^b_N = 0$ if the number of parallel interactions they use satisfies: $N > max\big(\frac{E[U^s]}{c}, \frac{E[U^b_N]+C(2)-2C(1)}{C(2)-C(1)}\big)$.*

Proof: From Lemma 2, we ascertain that in this case the seller agents accept any partnership, thus $(1 - F^s(x^s)) = 1$. Now notice that $C(1) + (C(2) - C(1))(N - 1) \leq C(N)$ for any $N > 1$ (because the cost function $C(N)$ is an increasing linear (or convex) function). Therefore if we use N that satisfies $E(U^b_N) < C(1) + (C(2) - C(1))(N - 1)$ (equivalent to $N > \frac{E[U^b_N]+C(2)-2C(1)}{C(2)-C(1)}$) then we can apply the same methodology used in Lemma 2 and prove that any buyer's expected utility when using this N value is negative and therefore the buyer agents use $x^b_N = 0$ as their reservation value. □

Returning to the proof of Theorem 4, notice that for any N that satisfies the theorem's condition, all agents' strategy would be to use a zero reservation value. Therefore, buyer and seller agents' expected utility becomes $V^b(0) = E[U^b_N] - C(N)$ and $V^s(0) = E[U^s] - Nc$, respectively (derived from Equation 4 and its appropriate modification for the seller as given in Appendix B). Since the buyer's utility $V^b(0) = E[U^b_N] - C(N)$ is negative for this N value (see Lemma 3), it is also a decreasing function of N for any N value that satisfies the theorem's condition. Therefore, a buyer agent will always have an incentive to deviate from strategy $(N', x^b_{N'} = 0)$ to strategy $(N, x^b_N = 0)$ for any $N' > N$. Thus there is no equilibrium in which buyer agents use a strategy $(N' > N, x^b_{N'})$. Since both $E[U^s] - c$ and $\frac{E[U^b_N]+C(2)-2C(1)}{C(2)-C(1)}$ are increasing concave functions (see Lemma 1), an N value that satisfies Equation 13 always exists. □

Once we have bounded the range of possible values for the equilibrium number of parallel interactions to be used by the buyer agents, we can suggest Algorithm 3 as a computational means for finding the equilibrium strategies.

Algorithm 3 An algorithm for finding the equilibrium strategies.

Input: ρ - required precision level.
Output: $Result$ - set of equilibrium strategy sets, in which each element defines the number of parallel interactions used by buyer agents and the reservation values used by the buyer and seller agents.

1: Set an upper bound, N_{upper} according to Theorem 4.
2: Set $Result = \emptyset$
3: **for** $i = 1$ to N_{upper} **do**
4:     Find the "suspected" equilibrium strategy sets using Algorithm 2 when the buyer agents use i parallel interactions.
5:     Add to $Result$ each of the sets found in step 4 that satisfies the conditions given in Theorem 3.
6: **end for**
7: **return** $Result$

While Algorithm 3 returns all the "suspected" equilibrium sets of strategies, we do not attempt to determine which one will eventually be used. The research on multiple non-dominating equilibrium strategies in game and agents theory is quite rich [17], and thus we do not include this question within the scope of the current paper.

Figure 6, illustrates the expected utility of buyer agents as a function of the number of parallel interactions, N, that they use when in environment 1. In this case, the optimal expected utility will be obtained when using $N = 11$. This expected utility that the buyer agents obtain in equilibrium is greater than the one obtained in the purely sequential two-sided search model (see Figure 1). The utility of the seller agents significantly drops in this case.

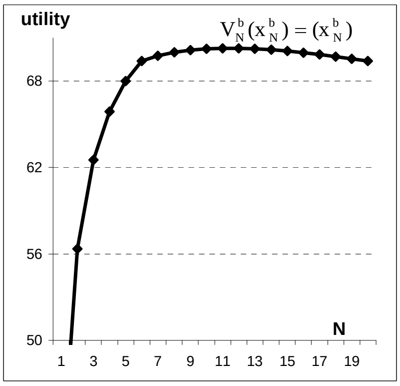

Figure 6: Extracting equilibrium sets of strategies for environment 1

3.5 Buyer Agents' Incentive to Use Parallel Search

As demonstrated in Figure 2, in many environments buyer agents have a strong incentive to adopt the new parallel search mechanism. The following Theorem 5 suggests a simple rule for checking whether or not buyer agents indeed have an incentive to adopt parallel search in a given environment.

Theorem 5 *If the index I, calculated according to Equation 17 below, is greater than 1 then any single buyer agent has an incentive to deviate from a purely sequential search strategy to a parallel one.*

$$I = \frac{C(1)\int_{y=x_1^b}^{\infty}(1-F^b(y)^2)dy}{C(2)\int_{y=x_1^b}^{\infty}(1-F^b(y))dy} \tag{17}$$

Proof: From Theorem 3 we learn that if a buyer agent's expected utility decreases when it shifts from its $(N=1, x_1^b)$ strategy to a new strategy $(N=2, x_2^b)$, while seller agents do not change their reservation values, then its expected utility when using strategy $(N=1, x_1^b)$ is greater than when using any other strategy $(N=k, x_k^b)$, $k>1$. Therefore, the adoption of a parallel search strategy occurs only if $V_2^b(x_2^b) > V_1^b(x_1^b)$ holds (while the seller agents use a reservation value x_1^s). Thus, given Theorem 1, all we need to prove is that if $I>1$ then $x_2^b > x_1^b$ holds.

Using two instances of Equation 9, one for $(N=1, x_1^b)$ and the other for $(N=2, x_2^b)$, while the seller type agents continue to use their equilibrium strategy x_1^s, we obtain:

$$1 = \frac{C(1)\int_{y=x_2^b}^{\infty}(1-F^b(y)^2)dy}{C(2)\int_{y=x_1^b}^{\infty}(1-F^b(y))dy} \tag{18}$$

If $x_2^b > x_1^b$ then replacing x_2^b with x_1^b in Equation 18 above will result in the right hand side of the equation being greater than 1. $\square$

Theorem 5 is applicable for any distribution function. Its advantage is that it supplies a simple tool for determining whether the analysis suggested in this section should be applied in the specific environment in which the agents operate. Using Equation 17, the agents (or a market maker or a market designer) can identify if a single buyer agent will have an incentive to deviate from the traditional sequential search strategy to a parallel search strategy (assuming all other agents use sequential search). If such an incentive exists, then a sequential-search based equilibrium will not hold in the environment and all buyer agents will use the parallel search technique. The new equilibrium in this case can be found using the algorithms developed throughout this section.

4 Future Extensions

The work presented in this paper is the first step towards adding parallel search into MAS two-sided search models. Obviously, once the analytical foundations for analyzing a model of this type are set, many further two-sided search variants can be considered. In this section we review two of these models, which we perceive as a natural extensions of our work. While the analysis of these two model variants is beyond the scope of this paper, we do choose to present the basic equations on which their analysis should be built (which are modifications of those given in the previous sections) and point to issues that need to be addressed as part of the attempt to apply the methodology presented above.

4.1 Considering Several Sellers at Each Search Stage

In the previous section we assumed that a buyer agent engaged in parallel interactions is required to reply to all the agents it interacts with simultaneously, thus the agent commits only to the seller associated with the highest utility at that search stage (assuming the utility is above its reservation value). Notice that within any search round of the buyer agent, there might be several seller agents, with whom it interacts, associated with a utility greater than its reservation value. In such a scenario, the buyer agent can benefit from an a-synchronous decision making mechanism, in which it holds its reply to some of the agents. This way, the buyer agent, upon receiving a rejection from the seller agent associated with the highest utility, can improve its expected utility by considering committing also to partnerships with other seller agents within the current search stage.

In this scenario, the buyer agent first identifies set A of seller agents to whom it is willing to commit (i.e., associated with a utility greater than its reservation value) out of those reviewed in the current search stage, and rejects the rest. Then it commits to the agent $A_j \in A$ that is associated with the partnership yielding the highest utility. If a rejection was received from agent A_j then this agent is removed from A and the buyer commits to a new partnership according to the same criteria. The process continues until either: (a) set A is empty, in which case the buyer agent initiates another search round; or (b) a dual commitment is obtained, in which case the agent rejects all remaining agents in A.

Notice that the above change does not affect the interaction protocol followed by the seller agents. Buyer agents, on the other hand, alter their interaction protocol (in comparison to the one presented in section 2.3) as follows:

loop
    Set a value N
    Locate randomly N seller type agents and initiate interaction with all of them in parallel
    Evaluate the utility $U^b_{(i)}$ from a partnership (transaction) with each agent i, ($i \leq N$).

Set $A = \{$all partnerships the agent is willing to commit to $\}$
Reject all partnerships not included in set A
while $(A \neq \emptyset)$ **do**
  Commit to the partnership associated with the highest utility in A and remove it from set A
  if (the other agent responded "commit") **then**
    Reject all the remaining agents in A
    Terminate search
  end if
end while
end loop

The analysis of this new mechanism variant requires the introduction of several additional definitions and notations to aid in formally introducing its equations. Consider a buyer agent that interacts with N seller agents at a specific search stage. We use $G_k(x)$ to denote the probability that the k^{th} best seller agent (in terms of associated utility) among the agents with whom the buyer agent interacts will offer a utility equal to or smaller than x. The value of $G_k(x)$ can be calculated using:[17]

$$G_k(x) = \sum_{i=1}^{k} \binom{N}{i-1} (F^b(x))^{N-i+1}(1 - F^b(x))^{i-1}; \quad x \geq 0 \tag{19}$$

Notice that the function $G_k(x)$ satisfies:

$$\lim_{x\to\infty} G_k(x) = 1; \quad \lim_{x\to 0} G_k(x) = 0 \tag{20}$$

The above is intuitive since the probability that the k^{th} best partnership will be associated with a utility smaller or equal to the maximum possible value is always 1. Similarly, the probability that the k^{th} best partnership will be associated with a utility smaller than the discrete value 0 is equal to zero (since the distribution is continuous). Notice that $G_k(x)$ increases as a function of k. This, also, is intuitive since the probability that the k-th best element in a sample of $N > k$ values is smaller than a value x is always smaller than the probability that the $(k+1)$-th best element in the sample is smaller than x.

Since $G_k(x)$ per se is a c.d.f. (commutative distribution function), we can find its associated p.d.f., denoted $g_k(x)$, by deriving Equation 19:

$$g_k(x) = \frac{dG_k(x)}{dx} = \sum_{i=1}^{k} \binom{N}{i-1} f^b(x)(F^b(x))^{N-i}(1 - F^b(x))^{i-2}(N-i+1-NF^b(x)) \tag{21}$$

Using the above definitions, we can calculate the expected value of a buyer agent when interacting with N seller agents while using a reservation value x^b and while the seller agents use a reservation value x^s. This is given by:

[17] The probability that the k^{th} best seller will have a utility smaller than x is the probability of having at least $N-k+1$ sellers with a utility equal to or smaller than x. Therefore the index i goes from 1 to k, adding the probability of having exactly $i-1$ sellers with a utility greater than x and $N-i+1$ sellers with a utility smaller or equal to x.

$$V^b(x^b) = \frac{\sum_{i=1}^{N} \Big(\overbrace{(1 - F^s(x^s))}^{\text{seller accepts buyer}} \ \overbrace{(F^s(x^s))^{i-1}}^{\text{former } (i-1) \text{ rejected}} \ \overbrace{\int_{y=x^b}^{\infty} y g_i(y) dy}^{\text{using the } i^{th} \text{ best seller}} \Big) - C(N)}{\sum_{i=1}^{N} (1 - F^s(x^s))(F^s(x^s))^{i-1}(1 - G_i(x^b))} \tag{22}$$

Here, the buyer agent would have liked to partner with the k^{th} best seller in the sample only if all $N - k$ better seller agents have turned it down and this k^{th} seller agent is associated with a better utility than the buyer agent's reservation value.

Similarly, the expected utility of a seller agent using a reservation value x^s while the buyers use a reservation value x^b and sampling N sellers at a time, and the other sellers use a reservation value $x^{s'}$ is given by:

$$V^s(x^s) = \frac{\Big(\int_{y=x^b}^{N} f^b(y) \sum_{i=1}^{N} \binom{N}{N-i} (F^b(y))^{N-i} (1 - F^b(y))^{i-1} (F^s(x^{s'}))^{i-1}\Big) \int_{y=x^s}^{\infty} y f^s(y) dy - c}{\Big(\int_{y=x^b}^{N} f^b(y) \sum_{i=1}^{N} \binom{N}{N-i} (F^b(y))^{N-i} (1 - F^b(y))^{i-1} (F^s(x^{s'}))^{i-1}\Big)(1 - F^s(x^s))} \tag{23}$$

Here, the term $\Big(\int_{y=x^b}^{N} f^b(y) \sum_{i=1}^{N} \binom{N}{N-i} (F^b(y))^{N-i} (1 - F^b(y))^{i-1} (F^s(x^{s'}))^{i-1}\Big)$ represents the probability that the current buyer agent will accept the specific seller agent for which the calculation takes place. This will happen only if the utility associated with this seller agent is above the buyer agent's reservation value, and as long as all other partnerships considered by the buyer agent, that are associated with a utility greater than the one of a partnership with the seller agent, have been rejected by the other sellers.

It is notable that in this model variant the seller agents' strategy is affected also by the strategy used by other seller agents in the environment. This, is in comparison to the model where buyer agents commit only to the best seller agent with whom they interact (if their perceived utilities are above their reservation value), in which seller agents' strategies are affected only by the buyer agents' strategy (see Equation 33 in Appendix B).

The optimal number of parallel interactions, N, in this model might be different from the one in the model presented in the previous sections. This is obviously a better model for the buyer type agents since it reduces their overall search costs. It is also beneficial for the seller type agents since it increases the probability of being accepted by the buyer, even if the seller is not associated with the highest utility in the buyer's sample. Nevertheless, since the interaction mechanism is required to be a-synchronous, various aspects need to be considered prior to the formal equilibrium analysis. For example, a seller agent in this model is required to wait until the buyer agent receives responses from several other seller agents. It is possible that during this time period this seller agent will be addressed by a new buyer with possibly a better utility associated with a partnership with it. Furthermore, the asynchronous communication requires defining appropriate protocols for maintaining timeouts and failures while communication takes place. We do believe however, that once these setbacks are resolved, the analysis of this model variant will benefit from the analysis methodologies presented in the previous section.

4.2 Adding the Parallel Search Technology to the Seller's Search

As noted, a model where sellers are also capable of conducting parallel search is not used in current markets. However, we do anticipate its application in future C2C marketplaces, where sellers will use more proactive methods to approach buyers. Though this is not the main focus of this paper, we provide a short review of such a model.

Using $C_s(N_s)$ and $C_b(N_b)$ to denote the seller and buyer agents' search costs, respectively, (where N_s and N_b are the sample sizes used by the seller and the buyer agents), we obtain:[18]

$$V^s(x^s) = E\Big[U^s_{best} \bullet 1[(U^b \geq U^b_{(j)} \forall j \leq N) \cap (U^b \geq x^b) \cap (U^s_{best} \geq x^s)] + \\ + V^s(x^s) \bullet 1[(\exists U^b_{(j)} > U^b, j \leq N) \cup (U^b < x^b) \cup (U^s_{best} < x^s)] - C_s(N_s)\Big] \quad (24)$$

resulting in:[19]

$$V^s(x^s) = \frac{(1 - F^b_{N_b}(x^b)) \int_{y=x^s}^{\infty} y f^s_{N_s}(y)dy - N_b C_s(N_s)}{(1 - F^s_{N_s}(x^s))(1 - F^b_{N_b}(x^b))} \quad (25)$$

Here, again, we obtain:

$$\lim_{x^s \to \infty} V^s(x^s) = \lim_{x^b \to \infty} V^b(x^b) = -\infty \quad (26)$$

and:

$$\lim_{x^s \to 0} V^s(x^s) = E(U^s_{N_s}) - \frac{N_b C_s(N_s)}{1 - F^b_N(x^b)} \quad (27)$$

We can also suggest a theory similar to Theorem 2, where the modification of Equation 9 is:

$$N_b C_s(N_s) = (1 - F_N(x))(E[U^s_{N_s}] - \int_{y=0}^{x^s} (1 - F^b_N(y))dy) \quad (28)$$

In this general model the equilibrium dynamics are highly influenced by the different cost function structures. The equilibrium in this model is a set of strategies defining both the number of parallel interactions and the reservation values that need to be used, ($x^s_{N_s}, x^b_{N_b}, N_s, N_b$), for each agent type. However finding such an equilibrium is not trivial (we need to guarantee that no single agent of any of the types will have an incentive to change the number of parallel interactions and/or the reservation value specified for its type) and we cannot guarantee that one actually exists.

The table in Figure 7 is an example of the reservation values that will be used by buyer and seller agents, when both agents use parallel search. It uses the same environment characteristics that are detailed in environment 1. Notice that each pair $(x^b_{N_b}, x^s_{N_s})$ is an equilibrium candidate (i.e., reflecting the stable set of reservation values used when the seller type agents use N_s parallel interactions and the buyer type agents use N_b). Therefore, each pair should be evaluated according to possibilities of using other N values and different reservation values. This is

[18] In this scenario, both seller and buyer equations are equivalent.

[19] Here we used: $(U^b_i > U^b_{(j)} \forall i \neq j, j \leq N) = \int_{x^b_{N_b}}^{\infty} F^b(y)^{N-1} f(y)dy = \frac{1 - F^b(x^b_{N_b})^{N_b}}{N_b}$.

		Seller Agent							
		1	2	3	4	5	6	7	8
Buyer Agent	1	(40.4 ; 40.4)	(24.2 ; 56.4)	(14.9 ; 62.5)	(8.4 ; 65.9)	(3.5 ; 68)	(-0.3 ; 69.4)	(-2.5 ; 69.8)	(-4.4 ; 70)
	2	(56.4 ; 24.2)	(43.1 ; 43.1)	(35.5 ; 50.2)	(30 ; 53.9)	(25.7 ; 56.1)	(22 ; 57.5)	(18.8 ; 58.5)	(15.9 ; 59.1)
	3	(62.5 ; 14.9)	(50.2 ; 35.5)	(43 ; 43)	(37.7 ; 46.9)	(33.4 ; 49.1)	(29.6 ; 50.4)	(26.2 ; 51.2)	(22.9 ; 51.7)
	4	(65.9 ; 8.4)	(53.9 ; 30)	(46.9 ; 37.7)	(41.5 ; 41.5)	(37 ; 43.6)	(33 ; 44.8)	(29.2 ; 45.3)	(25.6 ; 45.4)
	5	(68 ; 3.5)	(56.1 ; 25.7)	(49.1 ; 33.4)	(43.6 ; 37)	(38.9 ; 38.9)	(34.5 ; 39.7)	(30.3 ; 39.9)	(26.2 ; 39.6)
	6	(69.[illegible] ; -0.3)	(57.5 ; 22)	(50.4 ; 29.6)	(44.8 ; 33)	(39.7 ; 34.5)	(35 ; 35)	(30.3 ; 34.7)	(25.5 ; 33.8)
	7	(69.8 ; [illegible]2.5)	(58.5 ; 18.8)	(51.2 ; 26.2)	(45.3 ; 29.2)	(39.9 ; 30.3)	(34.7 ; 30.3)	(29.4 ; 29.4)	(24.1 ; 28)
	8	(70 ; -4.4)	(59.1 ; 15.9)	(51.7 ; 22.9)	(45.4 ; 25.6)	(39.6 ; 26.2)	(33.8 ; 25.5)	(28 ; 24.1)	(22 ; 22)

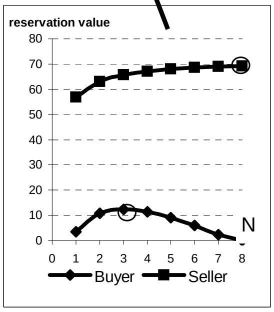

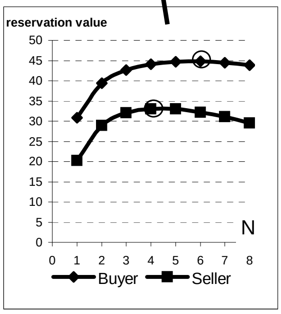

Figure 7: Dual parallel search - the upper value in each cell in the table is the buyer agents' utility and the lower value is the seller agents' utility, for each combination of parallel interactions (N_s, N_b) they use.

demonstrated by the two sets of graphs in Figure 7. In this example we learn that $(N_b = 4, N_s = 6)$ is an actual equilibrium, because neither agent has an incentive to use a different strategy. Contrarily, the pair $(N_b = 1, N_s = 5)$, is not an equilibrium point, since both agents can improve their expected utility by using different sample sizes ($N_b = 3$ for the buyer and $N_s = 8$ for the seller).

While this model suggests many interesting dynamics when analyzing the agents' strategies it is beyond the scope of this paper.

5 Related Work

The process of searching for partners, often associated with the agent matchmaking concept, has wide evidence in MAS literature [22, 7]. Service matchmaking and brokering have also been referred to in several systems and applications [43, 23]. In its wider context the pairwise partnership formation can be seen as part of the multi-agent coalition formation model found mostly in the electronic marketplace [19, 24, 36]. The emphasis on coalition formation mechanisms in ecommerce is usually associated with the advantage of potentially obtaining discounts based on volume as an incentive for buyer agents to cooperate [45, 46]. Additional coalition formation models for the electronic marketplace consider extensions of the transaction-oriented coalitions into long-term ones [7], and for large-scale electronic markets [24].

Nevertheless, while agents' search is recognized to be costly [9, 21], coalition and partnership formation models typically ignore this important factor. For example, in many mechanisms an assumption is made that agents can scan as many agents as needed [19], or that a central matcher or middle agents are available to support the matching process [10]. This applies also to coalition formation models where the analysis is derived from equilibrium considerations [5, 44]. Furthermore, many of these models do not consider the temporal continuity of the agents (i.e., their long term considerations).[20] Only a few works have considered the problem of finding matches for cooperative tasks without the help of a predefined organization or a central facilitator [31]. None of these has considered the parallel search option.

The concept of matching agents with other agents is an important function of many domains and processes (other than MAS), thus relevant research can also be found in economics and social studies (e.g., students applying to colleges, workers seeking employers, the marriage market, etc.). The review of economic literature reveals two additional relevant research areas. The first concerns the legacy two-sided matching application [14]. Here we find two types of agents where an agent of one type can be matched only with agents of the other type [34, 28]. A matching in this context is considered stable only if no two agents can be found who would prefer to be matched with each other rather than with their current matches. A special property of such two-sided markets is that stable matchings always exist and that the set of stable matchings is a lattice under the common preferences of all agents of the same type. These results were further used for the analysis of many extensions of the model [38, 20] and for empirical studies of two-sided search markets (e.g., the labor markets [33]). The main difference between the stable matching model and our decentralized costly two-sided search model is that the first considers a centralized match-making design, where the matching process does not involve costly search (i.e., the "stable matching" results are valid only if the process allows each searcher to consider all potential opportunities). In our two-sided search based model, this assumption that all parties involved have complete knowledge about the available options from which to choose is relaxed.

The second relevant economic research area, known as search theory, considers the problem of an individual interested in locating an opportunity which will minimize its expected cost (or maximize its expected utility), while the search process is associated with a search cost ([26, 25], and references therein). Within the framework of search theory, three main clusters of search models can be found. These are (a) the fixed sample size model; (b) the sequential search model; and (c) the variable sample size model. In the fixed sample size model, the searcher executes a single search round in which it obtains a large set of opportunities simultaneously [42] and chooses the one associated with the highest utility. In the sequential search strategy [35, 25] the searcher obtains a single opportunity at a time, allowing multiple search stages. Several attempts were made to adopt the fixed sample size search [21] and the sequential search [9] models in agent-based electronic trading environments associated with search costs. In these cases the main focus was on establishing the appropriate characteristics of the environment and search strategy rather than the computational aspects of extracting it. Last, the variable sample size search method [4, 13, 29, 30] suggests a combined approach in which several opportunities are obtained during each search period. Nevertheless, all these economic search models focus on establishing optimal strategies for the searcher, assuming no mutual search activities are taken. The environment in which the searcher operates in these models is completely static,

[20] In our model the agent has a temporal continuity, i.e., it is a continuously running process and not a "one-shot" computation that maps a single input to a single output and then terminates [11].

and is not affected by the searcher's strategy. Therefore, the searcher mainly needs to extract the strategy that optimizes its expected utility and does not need to model the effects changes in its strategy will have on the strategy of others in the environment. In contrast, the analysis of two-sided search is driven mainly by equilibrium considerations.

In an effort to understand the effect of dual search activities in such models, the "two-sided" search research followed. This notion was explored within the equilibrium search framework [1, 8, 40, 3, 27]. The main drawback of economic two-sided search models in the context of MAS environments is that they all consider a purely sequential search, sampling one partnering opportunity at a time, while agents' technology allows us to maintain a large number of parallel interactions when searching. Furthermore, while traditional two-sided search research is more concerned with describing the equilibrium equations and the equilibrium affect on global economic phenomena, MAS applications research requires computational means for deriving the agents' policies for different settings and the distributed computation of the equilibrium strategy.

6 Discussion and Conclusions

The parallel interaction technique is inherent in the infrastructure of autonomous information agents. Nowadays, as agents technology is a reality and traditional processing and communication limitations have been removed, it is high time to consider parallel search models in MAS domains and in particular the two-sided search application. Furthermore, as proven and demonstrated throughout the paper, there is a strong incentive for individual agents to use such techniques in many different environment settings.

We show that the agents can control their search intensity by initiating parallel interactions with other agents, for improving their utility. Furthermore, the mechanism we introduce is general, and can be applied to traditional search theory domains. For example, even in the classical marriage market application, nowadays, we can find equivalents for parallel search such as TV shows where a candidate becomes acquainted with several potential partners or even speed-dating.

We emphasize that the agent's utility will never decrease when using our proposed mechanism. This is because the agent controls the number of parallel interactions it uses in each search round, hence in the worst case scenario the proposed calculations will indicate that the optimal number of interactions is 1, thus the expected utility will be identical to the case where the traditional pure sequential method is used. In fact, the latter method is actually a specific case of our general model, using a single interaction over each search round.

Deriving the equilibrium strategy for the agents is a complex task. The novelty of the proposed analysis is that it results in efficient algorithms that allow the agents to quickly eliminate non-equilibrium strategies. Furthermore, we have managed to establish a simple thumb rule for determining in which environments the use of more than a single interaction at a time is beneficial.

While the model uses pairwise partnerships, we see a potential for extending it to limited size coalitions, e.g., a carpooling application in which the coalition size is limited to 4-5 members (based on the capacity of different car models). In this case, the set of equations used can be replicated to include a limited number of additional extensions of the coalition and the use of different opportunities distribution functions for each coalition size.

This paper reveals two potential future extensions of the analysis. The first concerns an alternative a-synchronous communication protocol the buyer agents can use to improve their performance and the second is a model in which all agents use the parallel search. Our initial results relating to these two extensions suggest that the analysis methodology and results given in this paper can nurture and play a significant role in the analysis and development of these interesting variants.

References

[1] A. Atakan. Assortative matching with explicit search costs. *Econometrica*, 74(3):667–680, 2006.

[2] Y. Bakos. Reducing buyer search costs: Implications for electronic marketplaces. *Management Science*, 42:1676–92, 1997.

[3] G. Becker. A theory of marriage. *Journal of Political Economy*, 81:813–846, 1973.

[4] J. Benhabib and C. Bull. Job search: The choice of intensity. *Journal of Political Economy*, 91:747–764, 1983.

[5] B. Bernheim, B. Peleg, and M. Whinston. Coalition-proof Nash equilibria: I concepts. *Journal of Economic Theory*, 42:1–12, 1987.

[6] S. Biswas and Y. Narahari. Analysis of supplier competition in electronic marketplaces. In *Proceedings of the International Conference on Automation, Energy, and Information Technology, EAIT-2001*, Indian Institute of Technology, Kharagpur, December 2001.

[7] S. Breban and J. Vassileva. Long-term coalitions for the electronic marketplace. In *B. Spencer, ed., Proceedings of E-Commerce Applications Workshop*, 2001.

[8] K. Burdett and R. Wright. Two-sided search with nontransferable utility. *Review of Economic Dynamics*, 1:220–245, 1998.

[9] S. Choi and J. Liu. Optimal time-constrained trading strategies for autonomous agents. In *Proc. of MAMA'2000*, 2000.

[10] K. Decker, K. Sycara, and M. Williamson. Middle-agents for the internet. In *Proc. of IJCAI*, 1997.

[11] O. Etzioni and D. Weld. Intelligent agents on the internet: Fact, fiction, and forecast. *IEEE Expert*, 10(3):44–49, 1995.

[12] P. Faratin, C. Sierra, and N. Jennings. Using similarity criteria to make issue tradeoffs in automated negotiations. *Artificial Intelligence*, 142(2):205–237, 2002.

[13] S. Gal, M. Landsberger, and B. Levykson. A compound strategy for search in the labor market. *International Economic Review*, 22(3):597–608, 1981.

[14] D. Gale and L. Shapley. College admissions and the stability of marriage. *American Math. Monthly*, 69:9–15, 1962.

[15] R. Guttman, A. Moukas, and P. Maes. Agent-mediated electronic commerce: A survey. *Knowledge Engineering Review*, 13(2):147–159, June 1998.

[16] M. Hadad and S. Kraus. Sharedplans in electronic commerce. In M. Klusch, editor, *Intelligent Information Agents*, pages 204–231. Springer Publisher, 1999.

[17] J. Harsanyi and R. Selten. *General theory of equilibrium selection in games*. MIT Press, Cambridge, Mass, 1988.

[18] M. He, N. Jennings, and H. Leung. On agent-mediated electronic commerce. *IEEE Transaction on Knowledge and Data Engineering*, 15(4):985–1003, 2003.

[19] T. Ito, H. Ochi, and T. Shintani. A group-buy protocol based on coalition formation for agent-mediated e-commerce. *International Journal of Computer and Information Science (IJCIS)*, 3(1):11–20, 2002.

[20] A. Kelso and V. Crawford. Job matching, coalition formation, and gross substitutes. *Econometrica*, 50(6):1483–1504, 1982.

[21] J. Kephart and A. Greenwald. Shopbot economics. *JAAMAS*, 5(3):255–287, 2002.

[22] M. Klusch. Agent-mediated trading: Intelligent agents and e-business. *Journal on Data and Knowledge Engineering*, 36(3), 2001.

[23] A. Kraft, S. Pitsch, and M. Vetter. Agent-driven online business in virtual communities. In *Proc. of HICSS-33*, Maui, Hawaii, January 2000.

[24] K. Lermann and O. Shehory. Coalition formation for large scale electronic markets. In *Proceedings of ICMAS'2000*, pages 216–222, Boston, 2000.

[25] S. Lippman and J. McCall. The economics of job search: A survey. *Economic Inquiry*, 14:155–189, 1976.

[26] J. McMillan and M. Rothschild. Search. In R. Aumann and S. Amsterdam, editors, *Handbook of Game Theory with Economic Applications*, pages 905–927. 1994.

[27] J. McNamara and E. Collins. The job search problem as an employer-candidate game. *Journal of Applied Probability*, 27(4):815–827, 1990.

[28] S. Mongell and A. Roth. Sorority rush as a two-sided matching mechanism. *American Economic Review*, 81(3):441–64, 1991.

[29] P. Morgan. Search and optimal sample size. *Review of Economic Studies*, 50(4):659–675, 1983.

[30] P. Morgan and R. Manning. Optimal search. *Econometrica*, 53(4):923–944, 1985.

[31] E. Ogston and S. Vassiliadis. Local distributed agent matchmaking. In *Proc. of the 9th International Conference on Cooperative Information Systems*, volume 2172, pages 67–79, 2001.

[32] S. Park, E. Durfee, and W. Birmingham. Use of markov chains to design an agent bidding strategy for continuous double auctions. *Journal of Artificial Intelligence Research (JAIR)*, 22:175–214, 2004.

[33] A. Roth. The evolution of the labor market for medical interns and residents: A case study in game theory. *Journal of Political Economy*, 92(6):991–1016, 1984.

[34] A. Roth and E. Peranson. The redesign of the matching market for american physicians: Some engineering aspects of economic design. *American Economic Review*, 89, 1999.

[35] M. Rothschild. Searching for the lowest price when the distribution of prices is unknown. *Journal of Political Economy*, 82:689–711, 1961.

[36] T. Sandholm, K. Larson, M. Andersson, O. Shehory, and F. Tohme. Coalition structure generation with worst case guarantees. *Artificial Intelligence Journal*, 111:209–238, 1999.

[37] 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.

[38] L. Shapley and M. Shubik. The assignment game i: the core. *International Journal of Game Theory*, 1:111–130, 1972.

[39] O. Shehory and S. Kraus. Methods for task allocation via agent coalition formation. *Artificial Intelligence Journal*, 101(1-2):165–200, 1998.

[40] R. Shimer and L. Smith. Assortative matching and search. *Econometrica*, 68(2):343–370, 2000.

[41] C. Sierra and F. Dignum. Agent-mediated electronic commerce: Scientific and technological roadmap. In F. Dignum and C. Sierra, editors, *Agent Mediated Electronic Commerce (AMEC2001)*, pages 1–18. Springer, 2001.

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

[43] K. Sycara, S. Widoff, M. Klusch, and J. Lu. Larks: Dynamic matchmaking among heterogeneous software agents in cyberspace. *Autonomous Agents and Multi-Agent Systems*, 5:173–203, 2002.

[44] F. Tohme and T. Sandholm. Coalition formation processes with belief revision among bounded rational self-interested agents. *Journal of Logic and Computation*, 9(6):793–815, 1999.

[45] N. Tsvetovat, K. Sycara, Y. Chen, and J. Ying. Customer coalitions in electronic markets. In *Proceedings of AMEC 2000*, pages 121–138, Barcelona, 2000.

[46] J. Yamamoto and K. Sycara. A stable and efficient buyer coalition formation scheme for e-marketplaces. In *Proceedings of the Fifth International Conference on Autonomous Agents*, pages 576–583, Montreal, Canada, 2001.

A The Buyer Agent's Expected Utility Function

Recall that the expected utility, $V^b(x^b)$, of a buyer agent that uses a strategy (N, x^b) while seller agents use a reservation value x^s is given by (Equation 1 from Section 3.1):

$$V^b(x^b) = E\Big[U^b_{best} \bullet 1[(U^b_{best} \geq x^b) \cap (U^s \geq x^s)] + V^b(x^b) \bullet 1[(U^b_{best} < x^b) \cup (U^s < x^s)] - C(N)\Big] \quad (29)$$

The term $\bullet 1[(U^b_{best} \geq x^b) \cap (U^s \geq x^s)]$ represents the indicator of the event where the specific buyer agent and its "best" encountered seller agent (in the current search round) have found the perceived utility from a transaction between the two of them to be greater than or equal to their reservation values, resulting in a dual commitment. In this case, the buyer agent's utility from the partnership is U^b_{best}, where U^b_{best} is characterized by the probability distribution function $f^b_N(x)$ (see Equation 3). Therefore, $E\Big[U^b_{best} \bullet 1[(U^b_{best} \geq x^b) \cap (U^s \geq x^s)]\Big] = (1 - F^s(x^s)) \int_{y=x^b}^{\infty} y f^b_N(y)dy$. The term $\bullet 1[(U^b_{best} < x^b) \cup (U^s < x^s)$ represents the indicator of the complementary event (i.e., where the seller and/or the buyer reject the potential partnership between them). The probability for the latter event is $1 - (1 - F^b_N(x^b))(1 - F^s(x^s))$, and the expected future utility (due to the stationarity of the problem) is $V^b(x^b)$.
Therefore, the above Equation 29 transforms into:

$$V^b(x^b) = (1 - F^s(x^s)) \int_{y=x^b}^{\infty} y f^b_N(y)dy + (1 - (1 - F^b_N(x^b))(1 - F^s(x^s)))V^b(x^b) - C(N) \quad (30)$$

resulting with Equation 4:

$$V^b(x^b) = \frac{(1 - F^s(x^s)) \int_{y=x^b}^{\infty} y f^b_N(y)dy - C(N)}{(1 - F^b_N(x^b))(1 - F^s(x^s))} \quad (31)$$

B Appropriate Seller Agents' Equation Modifications

In this appendix we present the appropriate modifications for seller agents for primary equations and theorems presented in the paper.

A modification of Equation 4

A seller agent's expected utility when using a reservation value x^s while buyer agents interact with N seller agents in parallel and use a reservation value x^b can be calculated using:

$$V^s(x^s) = \frac{\big(\int_{y=x^b}^{\infty} F^b_{N-1}(y) f^b(y)dy\big) \int_{y=x^s}^{\infty} y f^s(y)dy - c}{\big(\int_{y=x^b}^{\infty} F^b_{N-1}(y) f^b(y)dy\big)(1 - F^s(x^s))} \quad (32)$$

The term $\big(\int_{y=x^b}^{\infty} F^b_{N-1}(y) f^b(y)dy\big)$ represents the probability that the buyer agent commits to a partnership with the seller agent (i.e., the probability that the utility from partnering with the seller agent is greater than the maximum utility associated with partnering with any of the other $N-1$ seller agents with whom the buyer interacts). Recall that $\int_{y=x^b}^{\infty} F^b_{N-1}(y)f^b(y)dy = \int_{y=x^b}^{\infty} F^b(y)^{N-1}f^b(y)dy = (1-F^b(x^b)^N)/N$, thus:

$$V^s(x^s) = \frac{(1-F^b(x^b)^N)\int_{y=x^s}^{\infty} yf^s(y)dy - Nc}{(1-F^b(x^b)^N)(1-F^s(x^s))} \tag{33}$$

A modification of Equation 5

An immediate result from Equation 33:

$$\lim_{x^s\to\infty} V^s(x^s) = -\infty \quad ; \quad \lim_{x^s\to 0} V^s(x^s) = E[U^s] - \frac{Nc}{1-F^b(x^b)^N} \tag{34}$$

A modification of Theorem 1

Theorem 6 *The expected utility function $V^s(x^s)$ is quasi concave, with a unique maximum obtained at point x^s_N, satisfying:*

$$V^s(x^s_N){=}x^s_N \tag{35}$$

The proof is identical to the one used in Theorem 1.

An example where x^b_N cannot be computed directly from Equation 9

Consider for example the environment where utilities are normally distributed (a common distribution in most real-life applications). Here Equation 9 becomes:

$$C(N) = (1 - \int_{y=0}^{x^s_N} \frac{e^{\frac{-(y-\mu)^2}{2\sigma^2}}}{\sigma\sqrt{2\pi}}dy)(E[U^b_N] - \int_{y=0}^{x^b_N}(1 - \int \frac{e^{\frac{-(y-\mu)^2}{2\sigma^2}}}{\sigma\sqrt{2\pi}}dy)dy) \tag{36}$$

This function is certainly not trivial and contains non-integrated terms.

A modification for Theorem 2

Theorem 7 *Given the number of parallel interactions used by the buyer agents, N, and the reservation value they set, x^b_N, the seller agents' optimal reservation value, x^s_N, satisfies:*

$$cN = (1 - F^b_N(x^b_N))(E[U^s] - \int_{y=0}^{x^s_N}(1-F^s(y))dy) \tag{37}$$

The proof is identical to the one used in Theorem 2.

C Scenarios where the Expected Utility of an Agent Strictly Decreases as a Function of Its Reservation Value

As discussed in Proposition 1, there is a unique scenario, in which the expected utility function strictly decreases and never satisfies $V^b(x^b_N) = x^b_N$. The following theorem presents this scenario.

Theorem 8 *If the expected utility of a buyer agent when using a reservation value $x^b_N = 0$ (i.e., accepting any potential partnership) is negative, then the expected utility function $V^b(x)$ ($V^s(x)$ for the seller) strictly decreases.*

Proof: If $V^b(0) < 0$ then Equation 7 results in a negative first derivative at point $x^b = 0$. Therefore, a value x^b_N satisfying $V(x^b_N) = x^b_N$ is non-feasible (given the concavity characteristic found in Equation 8). Thus, given the limit found in Equation 5 for the case where $x^b \rightarrow \infty$, the expected utility function of the agent strictly decreases and reaches its maximum value (which is negative) for $x^b = 0$. □

Furthermore, it can easily be shown directly from Equation 4 that if $V^b(0) < 0$ then $V^b(x^b) < 0$ for any $x^b > 0$. Substituting $V^b(0) < 0$ in Equation 4 we obtain:

$$C(N) > \int_{y=0}^{\infty} y f^b_N(y) dy \tag{38}$$

Therefore, for any $x^b > 0$ the numerator in Equation 4 is always negative and the denominator is inevitably positive, resulting in a negative expected utility (thus a scenario where $V^b(x^b_N) = x^b_N$ does not exist).

Notation	Meaning
U^s , U^b	the seller and buyer agents' utility perceived from a specific transaction between them (respectively) (drawn from the probability distribution functions $f^s(U^s)$ and $f^b(U^b)$)
$f^s(U^s)$, $F^s(U^s)$	p.d.f and c.d.f functions of the utility obtained by seller agents from any potential transaction
$f^b(U^b)$, $F^b(U^b)$	p.d.f and c.d.f functions of the utility obtained by buyer agents from any potential transaction
$C(N)$	the buyer-agents' search cost when interacting with N seller agents (satisfying $\frac{dC(N)}{dN} > 0$ and $\frac{d^2C(N)}{dN^2} \geq 0$)
c	the seller-agents' search cost per search round (a seller agent interacts with a single buyer agent at a time)
S_{seller}: $U^s \rightarrow \{accept, reject\}$	seller agents' strategy (determines whether to accept or reject the current potential partnership that yields a utility U^s)
S_{buyer}: $U^b_{best} \rightarrow \{accept, N\}$	buyer agents' strategy (determines whether to accept or reject the best potential opportunity found in the current search stage (yielding a utility U^b_{best}) or to resume the search by interacting with N new seller agents)
U^b_{best}	the utility associated with the "best opportunity" among the potential opportunities available to the buyer agent at a given stage of its search
x^s	a reservation value used by a seller agent
x^b	a reservation value used by a buyer agent
$V^b(x)$, $V^s(x)$	the expected utility of buyer and seller agents, respectively, when using a reservation value x
$U^b_{(i)}$	the utility for a buyer agent obtained by a partnership with the $i-th$ seller agent with which it interacts at a specific search stage
x^b_N	the buyer agent's reservation value that maximizes its expected utility when maintaining N parallel interactions with seller agents at each search round, given the strategy used by the seller agents
x^s_N	the seller agent's reservation value that maximizes its expected utility when buyer agents maintain N parallel interactions with seller agents at each search round, given the strategy used by the buyer agents
$F^b_N(x)$	the cumulative distribution function of the maximum utility for the buyer agent when interacting with N seller agents in parallel
$f^b_N(x)$	the probability distribution function of the maximum utility for the buyer agent when interacting with N seller agents in parallel
$E[U^b_N]$	the expected maximum utility for the buyer agent when interacting with N seller agents in parallel
x^b_{N+} , x^b_{N-}	upper and lower bounds for the buyer agents' optimal reservation value, respectively
ρ	a precision level used in the suggested approximation means
I	an index used for determining if a single buyer agent has an incentive to deviate from a purely sequential search strategy to a parallel one.
A	a set of seller agents with whom a buyer agent is considering committing to a partnership, when trying to improve its acceptance probability (see section 4.1)
$G_k(x)$	the probability that the k^{th} best seller agent (in terms of associated utility) among the agents with whom the buyer agent interacts is offering a utility equal to or smaller than x (see section 4.1)
$g_k(x)$	the probability distribution function derived from the function $G_k(x)$ (see section 4.1)
$x^{s'}$	the reservation value used by other sellers, used in the process of setting a single seller agent's strategy in the variant where buyer agents are trying to improve their acceptance probability (see section 4.1)
$C_s(N_s)$, $C_b(N_b)$	the seller and buyer agents' search costs, respectively, when all agents use parallel search (see section 4.2)
N_s , N_b	the sample sizes used by the seller and the buyer agents, respectively, when all agents use parallel search (see section 4.2)

Table 1: Summary of notations used for special variables and constants.