

Multi-Goal Economic Search using Dynamic Search Structures

David Sarne · Efrat Manisterski · Sarit Kraus

Received: date / Accepted: date

Abstract This paper investigates cooperative search strategies for agents engaged in costly search in a complex environment. Searching cooperatively, several search goals can be satisfied within a single search effort. Given the searchers' preferences, the goal is to conduct a search in a way that the expected overall utility out of the set of opportunities found (e.g., products when operating in a market) minus the costs associated with finding that set is maximized. This search scheme, given in the context of a group search, applies also to scenarios where a single agent has to search for a set of items for satisfying several different goals. The uniqueness of the proposed mechanism is in the ability to partition the group of agents/goals into sub-groups where the search continues for each group autonomously. As we show throughout the paper, this strategy is favorable as it weakly dominates (i.e., can improve but never worsen) cooperative and autonomous search techniques. The paper presents a comprehensive analysis of the new search method and highlights the specific characteristics of the optimal search strategy. Furthermore, we introduce innovative algorithms for extracting the optimal search strategy in a range of common environments, that eliminates the computational overhead associated with the use of the partitioning technique.

1 Introduction

In many Multi-Agent based environments autonomous agents are required to engage in search in order to learn what opportunities and options are available to them. For example, a buyer agent in an electronic marketplace often needs to explore and identify the available opportunities to buy a specific requested product [25]. In most cases, the search performed for learning about new opportunities is costly. Search costs reflect the resources (not necessarily monetary) that need to be invested/consumed for maintaining search-related activities, such

This work was supported in part by NSF grant no IIS0705587, the Israeli Science Foundation grants no 1685 and 1401/09 and BSF grant 2008-404. Kraus is also affiliated with UMIACS. Preliminary parts of the analysis given in this paper appeared in [27].

Department of Computer Science,
Bar-Ilan University, Ramat Gan, 52900 Israel

as locating opportunities, analyzing and comparing them and negotiating over them with other agents. The existence of search costs in Multi-Agent Systems (MAS) and the need to take these into consideration when setting the agents' strategies is widely recognized [1, 9, 20]. The influence of these costs increases as a function of the amount and richness of opportunities that can be potentially found in the MAS environment.¹ Therefore, a key principle in an agent's search process (referred to, in this context, as "economic search"²) [24, 29], is to identify the point where the expected benefit from further exploration of the environment is smaller than the cost associated with the additional search effort.

An important form of "economic search" is a multi-goal-satisfaction search, i.e., when the information gathered along the search applies to several different goals that need to be satisfied. The goals can either be of a single searcher or several different searchers, and the idea is that instead of conducting a separate search for satisfying each goal, a coordinated search takes place in a way that all goals are eventually satisfied. For example, a single searcher may search for several different goods [14, 8, 7], e.g., presents for several of her friends, when visiting different department stores. This way, instead of searching for a present to one friend at a time (and upon buying a present restarting the search process for a present to the next friend) the searcher can shop for all presents at the same time, as many of the items sold may be of value as a present to more than one friend. Similarly, several agents, possibly with different sets of preferences, may decide to coordinate their search, or have one of them (or a designated new agent) search on their behalf [17, 28]. For example, consider a group of buyers interested in buying a new car of a specific brand and model. The same type of car can be found in many dealerships and obviously each buyer values differently the different terms and conditions offered by each dealership (e.g., warranty period, payment options, dealership's reputation). Nevertheless, in order to learn these terms and conditions offered by each dealership for the car one needs to physically visit the different dealerships. This process is associated with a substantial search cost, as it involves spending time (driving, querying sellers) and money (e.g., for fuel). Alternatively, we can think of a cooperative-based search, by which one buyer is searching on behalf of the group and the group compensates her for her search expenses or each buyer, in her turn, visits a different dealership and shares the information with all others (in a way that the overall search cost and effort are split among everyone). This does not necessarily imply that all buyers will buy the car eventually at the same dealership, however in this manner the group can share, reuse and re-allocate opportunities (that otherwise, if using separate searches, might have been discarded) among themselves.

When considering a multi-goal search one needs to take into consideration additional overhead costs. These costs can be of various forms. For example, in the single searcher's case, these costs can be the additional costs associated with the need to evaluate the value of each opportunity found to each goal (e.g., when searching for presents to few of her friends the searcher needs to consider, whenever introduced with a new product, the value of this item to each friend to whom it might be suitable as a gift). In the multi-agent cooperative-based search, these costs can be the communication, coordination and internal organization costs, both for the formation of the cooperation and for its maintenance. This overhead is usually modeled as an increasing function of the number of goals that need to be satisfied or the number of agents coordinating their search [39, 21, 38].

¹ Even in settings where technology can reduce the cost of exploring a single opportunity (e.g., when buying a specific product over the internet) the cost of evaluating all possible opportunities available becomes an important parameter affecting the agents' search strategy.

² "Economic search" can also be considered as sampling. This, as opposed to the classical AI search [16] in which an agent seeks a sequence of actions that will bring it from an initial state to a requested goal state.

Former analysis of multi-goal search [14,8,7,38] suggests a significant potential improvement for the agents' performance when using the multi-goal search method in various environments (in comparison to individual searches [29]). However this improvement cannot be generally guaranteed for all environments. This is principally because the strategies proposed for the multi-goal search assume that the set of goals is fixed throughout the entire search process. This assumption limits the search strategy to the decision of whether to continue the search as a whole or terminate the search completely. In this paper, we suggest integrating in the search strategy also the option to continue the search with only a subset of goals (while either satisfying the remaining goals with the findings obtained so far or continuing the search for satisfying these unsatisfied goals as a parallel effort). This form of search is more realistic and applicable in comparison to the models of fixed goal-sets. For example, a single searcher seeking for several different goods does not need to delay the purchase of all items until obtaining satisfying price quotes for all the different products she is interested in. Instead, she can buy some of the products when appealing opportunities to buy them are found, and avoid the overhead associated with searching for these items any further. Similarly, in the cooperative multi-agent search example, it might be beneficial for some of the agents to exploit some of the opportunities found so far and have the rest of the agents continue the search, cooperatively, with smaller overall overhead (e.g., upon recognizing a dealership that fits perfectly the needs of a subgroup of buyers, have this subgroup buy at that dealership, and the rest of the group continue the search with a significantly smaller overhead).

We show that the new search method generally guarantees at least the performance achieved by having each agent search individually and the one obtained by the cooperative search without the proposed enhancement. The main advantage of applying intelligent restructuring (i.e., partitioning) strategies along the search is that it takes into consideration the expected contribution of any subset of goals to the overall utility and to the cost of search at any search stage. This, as opposed to merely considering the gain in having the group maintain the search as a whole. The proposed cooperative search technique is particularly valuable in scenarios where different goals reflect non-correlated (or partially correlated) preferences. Here, we may identify opportunities in which it is more beneficial to terminate the search for some of the goals (taking advantage of any of the currently known opportunities) while resuming the search cooperatively in a reduced form (e.g., continue the search only for the remaining goals), hence with a smaller overhead.

While the analysis given in the paper is applicable both for the case of a single agent search, when attempting to satisfy multiple goals, and for the case where the multi-goal search is associated with a joint effort of several different agents, its formal presentation is given within the framework of the latter model. This is primarily due to the fact that the model of multi-goal search when performed cooperatively by a group of agents encompasses a richer formulation and can be used as an infrastructure for further research, investigating various important aspects of the process relating to the formation and maintenance of cooperation.

The contributions of this paper are twofold. First, we introduce and analyze a new cooperative search technique by which the search can be partitioned into several sub-efforts, each conducted separately for satisfying a subset of goals, as part of the search strategy space. The new search technique contributes directly both to multi-goal search and coalitional sampling models, as part of the economic search-theory domain, and to coalition formation and cooperation theories as developing the optimal cooperative search of a coalition lays the foundations for analyzing important questions such as the formation of cooperation and stability and payoff division mechanisms [34,45].

We use VSCS (Variable Structure Cooperative Search) to denote our new cooperative search method, as opposed to the Fixed-Structure Cooperative Search technique, denoted FSCS, by which the group continues as a whole throughout the search process [14, 8, 7, 38]. The individual search technique (i.e., where each agent searches separately for satisfying a single goal at a time is denoted SAS (Single Agent Search [29]). Alongside the formal introduction of the model, we supply an in-depth theoretical analysis leading to the specific characteristics of the optimal strategy to be used in the new search model. As we show throughout the paper, both the cooperative (FSCS) and individual (SAS) search techniques can be considered specific cases of our proposed search technique, VSCS (and are weakly dominated by VSCS in terms of performance).

While the VSCS model improves the agents' performance, it also significantly affects the computational complexity of computing the optimal search strategy. This is mainly because the strategy space in the new model increases immensely (due to the many partitioning alternatives). As we show in the following sections, when constructing the new search strategy, we are no longer concerned with only when to continue the search, but rather need to decide under which structure to continue. For this purpose we need to evaluate the potential benefit from any possible partition of the search as well as additional future partitions of any of the agent subsets throughout any future sequence of opportunities encountered. This imposes a significant computational challenge. Therefore, a second important contribution of this paper is the introduction of a computational algorithmic-based means for facilitating the calculation process of such a strategy in common MAS environments (e.g., environments characterized by fixed or non-increasing marginal search costs). The uniqueness of our algorithm relates to the fact that its computational complexity is similar to the one characterizing the FSCS model and does not depend on the potential number of partitions.

In the following section we address relevant literature in the area of MAS cooperation, coalition formation and economic search which apply to single-agent multi-goal as well as cooperative search. Then we formally introduce the VSCS model, present its analysis and introduce appropriate computational means as described above. Since the superiority of the new model over the two other methods is analytically straightforward, the focus of the paper remains on strategy analysis. Consequently, the different specific characteristics of the proposed model are illustrated using designated synthetic small-scale environments, designed for this purpose. Nevertheless, towards the end of the paper we supply a comparative example for the performance of the new cooperative search method, based on authentic settings, fully correlated with a genuine eCommerce specific vertical market. Our computational algorithm runs within less than a second in this setting which contains thousands of opportunities.

2 Related Work

The problem of a searcher operating in a costly environment, seeking to maximize his long term utility is addressed in classical economic search theory ([24, 29], and references therein). Among the three main search models that were introduced (the fixed sample size model [40], the sequential search model [35, 24] and the combined model [13, 32]) the most relevant for our model is the sequential search, where the searcher draws one observation at a time, allowing multiple search stages. Attempts to adopt the sequential search model for agent-based electronic trading environments associated with search costs are suggested in [9, 20]. The main focus of these works is on establishing the appropriate characteristics of the environment and search strategy rather than the computational aspects of extracting it.

Few studies in the area of economic search theory have attempted to investigate the problem of a single searcher, attempting to purchase several commodities (instead of one in the classical economic search models) while facing imperfect information concerning prices [14,8,7]. For example, Burdett and Malueg [7] extend the single search model and analyze a consumer visiting one or more stores in order to minimize the total expenditure. Carlson and McAfee in [8] extend Burdett and Malueg's work and analyze the case where customers have no perfect recall (i.e., the ability to exploit formerly explored opportunities). Gratti [14] identifies necessary and sufficient conditions in an agent's utility function that guarantee that the optimal search strategy can be characterized by the reservation property. Common to all these multi-goal-satisfaction models is that the set of goals to be satisfied throughout the search is stationary and does not change along time.

Substantial research efforts have been devoted over the years for studying cooperation among agents, as in many multi-agent environments autonomous agents may find benefit in cooperating and coordinating their actions. The cooperation is mostly useful when an agent is incapable of completing a task by itself or when operating as a group can improve the overall performance [5,22,15]. Consequently, group based cooperative behavior has been suggested in various domains [42,12,43,44]. The recognition of the advantages encapsulated in teamwork and cooperative behaviors, is the main driving force of many coalition formation models in the area of cooperative game theory and MAS [23,39]. A review of the extensive literature on coalition formation can be found in [18]. While coalition formation and coordination models can be widely found in the electronic market domain, most work emphasizes mechanisms for forming cooperation for the purpose of aggregating demands in order to obtain volume discount [43,44]. Additional coalition formation models for the electronic marketplace consider extensions of transaction-oriented coalitions into long-term ones [5], and for large-scale electronic markets [22]. Overall, the majority of cooperation and coalition formation MAS-related research tends to focus in the way coalitions are formed and consequently concerns issues such as the optimal division of agents into disjoint exhaustive coalitions [41,44], division of coalition payoffs [44] and enforcement methods for interaction protocols [30]. Only several authors have considered the problem of determining the strategy of a group once formed in the electronic commerce domain, once the coalition is formed [17], and no work to date considers strategies for a cooperative search in a costly environment where partitioning capabilities can be applied.

3 Cooperative Search

For exposition purposes, we adopt the legacy buyers-group application for the electronic marketplace [43,44] and in particular the B2C (Business-to-Consumer) market, where sellers can supply almost any volume of demand. Notwithstanding, we emphasize that the cooperative search strategy we present in this paper is general and can be applied to any scenario where a single searcher can satisfy several of its search goals as part the same search process or in the multi-agent case when the opportunities obtained are applicable to several agents.

Our environment description extends the multi-goal search [14,38,8] to support partitioning as part of the set of actions in the strategy definition. We consider an electronic marketplace where numerous buyer and seller agents can be found, each interested in buying or offering to sell a well defined product. A product can be offered by many different seller agents under various terms and policies (including price). We assume that while buyer agents are ignorant of individual seller agents' offers, they are acquainted with or can learn the overall distribution of opportunities (whereas an opportunity is defined as the option to

buy the desired product under specific terms and policies) in the marketplace. The latter assumption is standard and common in economic search models (e.g., [4, 13, 31, 32]).³

In the absence of central matching mechanisms or mediators, each agent needs to search for appropriate opportunities to buy its requested product. This process, in which a single agent searches autonomously for satisfying a single search goal (SAS), is described and analyzed in [29]. Throughout its search the buyer agent locates seller agents sequentially (i.e., one at a time) and learns about their offers by interacting with them. Upon learning a new opportunity details, the buyer agent evaluates it using its own utility function. Similar to many consumer theories [11], we apply the multi-attribute utility theory (MAUT) [19] when evaluating the different opportunities found. This enables a set of preferences to be represented by a numerical utility function. We consider the agents to be heterogeneous; each having its own utility function defined over product attributes, terms and policies as well as reputation and trust factors. Based on its evaluation the agent makes a decision whether to exploit any of the opportunities encountered until this point (i.e., buy from any of the sellers located) or to resume its search in a similar manner. The optimal search strategy in a single-agent's search is stationary (i.e., it does not change from one search stage to another): the agent sets a lower limit stopping rule and terminates the search when reaching an opportunity that yields a utility greater than or equal to this limit [29].

The search activity is assumed to be costly [9, 20, 28] in a way that each search stage of the process induces a search cost. We assume utilities and costs are commensurable and additive. Furthermore, as in other MAS cooperation models [18, 36], we assume that there is an option for side payments. Recognizing the benefits of a cooperative search, buyer agents, interested in similar products or interchangeable products, may decide to establish such a cooperative search effort. This process, which can be seen as a coalition formation process [41, 43] involves several important aspects, relating to the way the coalition is generated and the way it executes its cooperative search task and divides the generated value among the coalition members [41, 44]. The focus of this paper is in the cooperative search itself, i.e., finding an optimal cooperative search strategy for the coalition once it was formed, based on the utilities reported by the different coalition members. The optimal cooperative search strategy is the one that maximizes the overall utility of the coalition members as reflected by their (not necessarily truthful) reported utilities less the costs accumulated along the search process. This is equivalent to the problem of a single searcher in a multi-goal search as the goal is to maximize the accumulated utility obtained from satisfying the different goals while taking into account the cost of the search. The remaining coalition-formation-related aspects are beyond the scope of this paper and can be investigated separately. In fact, given the overall utility maximization goal, the coalition's optimal cooperative search strategy is not influenced by the coalition formation mechanisms but rather influences them. This is because the agents can only benefit (in terms of their reported utility) from any increase in the coalition's expected overall utility and thus the overall utility maximization strategy is the preferred strategy by all agents at every stage of the search (i.e., no conflict of interests as of how the cooperative search should be executed).⁴

³ There are several methods by which an agent can be acquainted with this distribution function. For example, agents can rely on recent history (e.g., using spectators) and past experience. This method can be useful even in dynamic on-line markets where despite frequent changes in individual opportunities, the overall distribution remains steady (e.g., price dispersion in many well-established online retail markets [2, 3, 6, 10]). Another alternative that can be used is Bayesian sampling which can be added to the model. Since this is not the main focus of the paper, we do not include these stochastic elements.

⁴ We leave stability issues out of the scope of this paper.

There are many ways cooperative search can take place, e.g., assigning a representative agent that searches on behalf of the coalition or simply taking turns searching, each deriving a different search cost overhead structure. For this paper's purposes, there is no importance to the method used, as long as the exploration of opportunities is performed sequentially. Once the agents decide to cooperate, and set the rules for the cooperation they can simply be treated as a coalition that encounters a single seller at each search stage, thus accumulating new opportunities along time. As in various other coalition formation models, we assume that the costs associated with maintaining the coalition are an increasing function of the size of the coalition [39,21,28] (or of the number of goals that need to be satisfied in the single searcher's case, as each opportunity needs to be evaluated in the context of every goal).

The cooperative search strategy determines the action to be taken after obtaining and reviewing a new opportunity for the coalition members (i.e., upon completion of each future search stage). While in the fixed structure model (FSCS) the set of possible actions for each stage includes: (a) terminating the search; and (b) resuming the search while keeping the current structure; in the new model (VSCS) the coalition can also choose to partition itself into sub-coalitions (where each member in the partitioned coalition is assigned to one of the sub-coalitions) that set their optimal search strategy from this point onward, independently (i.e., search autonomously or terminate their search).⁵

Before moving on to the analysis sections, and formally introducing the model, we recap the assumption used and correlate them to the different entities combining the model:

- **Agents/Goals** - agents can form and restructure cooperation, and individuals can combine and separate goals into different search efforts; learn about the benefit from different opportunities through a costly search.
- **Opportunities** - complex however can be valued by a numerical utility function; each potentially applies to all agents or satisfy all goals (and can be selected for as many agents/goals requested eventually); the distribution of opportunities in the environment is known or can be learned; can be evaluated only through costly search;
- **The Search Process** - assumed to be costly in a way that each search stage of the process induces a search cost, which depends on the number of goals or the number of cooperating agents; opportunities are explored sequentially, in a random order; the optimal cooperative search strategy is the one that maximizes the overall utility of the coalition members (or of all goals in a multi-goal search), while taking into account the cost of the search.
- **Utilities and Costs** - are commensurable and additive.

4 Finding the Coalition's Search Strategy

The following section formally defines the search environment and the coalition's search strategy. For convenience, all the notations given, and their meanings, are summarized in a table at the end of the paper.

Let $B = (B_1, B_2, \dots, B_{|B|})$ be the set of the attributes defining any of the potentially available opportunities in the market, where each attribute B_i can be assigned a value from the finite set $(b_{min}^i, \dots, b_{max}^i)$. An opportunity's type is defined by the vector $\vec{o}_i = (b_1, b_2, \dots, b_{|B|})$,

⁵ The suggested partitioning process can be seen as coalition reconstruction [33], and is very similar in its nature to the coalition formation process. It is used whenever the coalition can increase its overall utility by acting as several autonomous sub-coalitions. In fact, the partitioning principles can be determined and agreed as a condition for forming the coalition.

assigning a value b_i to each specific attribute B_i .⁶ Throughout the search, identical opportunities (i.e., of the same type) may be encountered. We use O to denote the space of potential opportunity types the coalition may encounter. The opportunity types' distribution in the marketplace is denoted by the probability function $p(\vec{o})$, $\sum_{\vec{o} \in O} p(\vec{o}) = 1$. We consider a coalition $A_g = \{a_1, a_2, \dots, a_{|A_g|}\}$, where a_j is the j -th buyer agent in the coalition. Each buyer agent, a_j , evaluates opportunities using a utility function $U_j : O \rightarrow R$, where $U_j(\mathbf{o})$ is the agent's utility from opportunity type \mathbf{o} . The search cost associated with having a coalition of size n (i.e., having n agents in the coalition) for each search round is denoted by the function $c(n)$.

We represent the world states of a coalition or a sub-coalition using the set of opportunities known to them. Given a set $\theta = \{\mathbf{o}_1, \dots, \mathbf{o}_l\}$ of opportunities known to coalition A_g (i.e., a set of l opportunities encountered at former search stages of coalition A_g), it is sufficient to maintain a subset s of θ to represent the current state of this coalition. Subset s stores the opportunities from θ that maximize the utility of each of the agents in A_g . Formally, we can calculate the state s of a coalition A_g acquainted with a set θ of known opportunities by using the function: $s = \text{state}(A_g, \theta) = \{\mathbf{o}_{a_j}^s | a_j \in A_g, \mathbf{o}_{a_j}^s \in \theta, U_j(\mathbf{o}_{a_j}^s) \geq U_j(\mathbf{o}), \forall \mathbf{o} \in \theta\}$.⁷ We use S_{A_g} to denote the set of all possible states of a coalition A_g . Reaching a state s , the expected utility of a coalition A_g when using its optimal strategy (from this point onwards) is denoted $V^*(A_g, s)$.

The following example illustrates the computation of the current state of different coalitions as a function of the information they have (i.e., the opportunities they are currently familiar with).

Example 1 Consider the following environment:

Environment 1 There are 3 agents $\{a_1, a_2, a_3\}$ evaluating the usefulness of a cooperative search in a market associated with 4 types of opportunities $\{\mathbf{o}_1, \mathbf{o}_2, \mathbf{o}_3, \mathbf{o}_4\}$, equally distributed ($p(\mathbf{o}_i) = 0.25, \forall i$). The agents' utilities associated with each opportunity type are presented in the following table:

Opportunities	a_1	a_2	a_3
\mathbf{o}_1	0.1	0.2	1
\mathbf{o}_2	0.3	1	0.2
\mathbf{o}_3	1	0.1	0.3
\mathbf{o}_4	1.3	1.1	2.3

Given a coalition of agents and a set of known opportunities the coalition's current state includes the opportunities that maximize the coalitions members' utilities. The following table describes the different states of the different coalitions given some of the different possible sets of opportunities encountered.

coalition \ θ	$\{\mathbf{o}_1, \mathbf{o}_2\}$	$\{\mathbf{o}_1, \mathbf{o}_2, \mathbf{o}_3\}$	$\{\mathbf{o}_1, \mathbf{o}_2, \mathbf{o}_3, \mathbf{o}_4\}$
$\{a_1, a_2, a_3\}$	$\{\mathbf{o}_1, \mathbf{o}_2\}$	$\{\mathbf{o}_1, \mathbf{o}_2, \mathbf{o}_3\}$	$\{\mathbf{o}_4\}$
$\{a_1, a_2\}$	$\{\mathbf{o}_2\}$	$\{\mathbf{o}_2, \mathbf{o}_3\}$	$\{\mathbf{o}_4\}$
$\{a_1, a_3\}$	$\{\mathbf{o}_1, \mathbf{o}_2\}$	$\{\mathbf{o}_1, \mathbf{o}_3\}$	$\{\mathbf{o}_4\}$
$\{a_2, a_3\}$	$\{\mathbf{o}_1, \mathbf{o}_2\}$	$\{\mathbf{o}_1, \mathbf{o}_2\}$	$\{\mathbf{o}_4\}$
$\{a_1\}$	$\{\mathbf{o}_2\}$	$\{\mathbf{o}_3\}$	$\{\mathbf{o}_4\}$
$\{a_2\}$	$\{\mathbf{o}_2\}$	$\{\mathbf{o}_2\}$	$\{\mathbf{o}_4\}$
$\{a_3\}$	$\{\mathbf{o}_1\}$	$\{\mathbf{o}_1\}$	$\{\mathbf{o}_4\}$

⁶ Notice that \mathbf{o} is noted as a vector since it assigns a specific value to each of the different attributes, terms and conditions associated with a specific opportunity. For example, a specific opportunity to buy a calculator can be represented by the vector $\mathbf{o} = (\text{scientific}, 20\$, \text{small Display}, \text{pocket}, 1\text{YR warranty})$.

⁷ If more than one maximizing opportunity exists, $\text{state}()$ will return a single opportunity according to a predefined order.

As can be seen from the example, in the proposed state presentation method only the effective opportunities in are stored in the state set, thus many redundant states are represented as a single state.

4.1 Coalition's Optimal Strategy for a General Search Cost Structure

We begin our analysis by constructing the coalition's expected utility as a function of the coalition's strategy along its future search. Consequently, we obtain the appropriate equations from which the optimal strategy can be extracted.

Consider the scenario where a coalition A_g chooses to terminate the search at state s . Here, the coalition's utility, denoted $V_r(A_g, s)$, is the aggregated coalition member's utilities when each coalition member, a_j , is assigned the opportunity $\mathbf{o}_{a_j}^s$ which maximizes its utility function, U_j , from the set of currently known opportunities in s :

$$V_r(A_g, s) = \sum_{a_j \in A_g} U_j(\mathbf{o}_{a_j}^s) \quad (1)$$

Next, we consider the scenarios in which coalition A_g resumes its search at state s while keeping its current structure. Here we can divide the opportunities space into two sub-spaces, $O_{improve}^s$ and O_{stay}^s based on the resulting state transition incurred if the coalition locates an opportunity of one of the types they contain. The sub-space $O_{improve}^s$ contains all opportunity types that change the coalition state (if added to the set θ)⁸ while the sub-space O_{stay}^s contains all the opportunity types that will not change the coalition's current state. More specifically we consider the two following cases:

(1) When the coalition encounters an opportunity \mathbf{o} that belongs to O_{stay}^s , the coalition's current state is still s , i.e., $state(A_g, s \cup \{\mathbf{o}\}) = s$. In this case the coalition's expected utility does not change due to the new opportunity. This derives from the stationary nature of the problem - if no better state is reached, the search resumes using the same strategy, yielding the same expected utility from this point onward.

(2) When the coalition encounters opportunities that belong to $O_{improve}^s$ the coalition's current state changes to $s' = state(A_g, s \cup \{\mathbf{o}\})$. Since we assume that the coalition will act according to the optimal strategy at any future state $s' \neq s$, the coalition's expected utility once it reaches the new state s' can be expressed as $V^*(A_g, s')$.

Using the above distinction, the expected utility of the coalition when all agents in A_g resume the search, denoted $V_r(A_g, s)$, is attained by:

$$V_r(A_g, s) = \sum_{\mathbf{o} \in O_{improve}^s} p(\mathbf{o}) V^*(A_g, state(A_g, s \cup \{\mathbf{o}\})) + \sum_{\mathbf{o} \in O_{stay}^s} p(\mathbf{o}) V_r(A_g, s) - c(|A_g|) \quad (2)$$

Applying some basic mathematic manipulations to the above equation, we obtain:

$$V_r(A_g, s) = \frac{\sum_{\mathbf{o} \in O_{improve}^s} p(\mathbf{o}) V^*(A_g, state(A_g, s \cup \{\mathbf{o}\})) - c(|A_g|)}{1 - \sum_{\mathbf{o} \in O_{stay}^s} p(\mathbf{o})} \quad (3)$$

Since $1 - \sum_{\mathbf{o} \in O_{stay}^s} p(\mathbf{o}) = \sum_{\mathbf{o} \in O_{improve}^s} p(\mathbf{o})$, we obtain:

⁸ Notice that the transition is always to an improved state since the coalition can never worsen its state when learning of a new opportunity.

$$V_r(A_g, s) = \frac{\sum_{\mathbf{o} \in O_{improve}^s} p(\mathbf{o}) V^*(A_g, state(A_g, s \cup \{\mathbf{o}\})) - c(|A_g|)}{\sum_{\mathbf{o} \in O_{improve}^s} p(\mathbf{o})} \quad (4)$$

The third possible scenario is where coalition $A_g = \{a_1, \dots, a_{|A_g|}\}$ partitions into a set $P = (A_1, \dots, A_k)$ of disjoint non-empty sub-coalitions ($A_i \cap A_j = \emptyset, \forall i, j \leq k, i \neq j, \bigcup_{i=1}^k A_i = A_g$) that set their search strategy independently. We denote by M_{A_g} the set of all possible partitions of coalition A_g . The selected partition will be the one yielding the maximum expected utility, assuming all the sub-coalitions created use their optimal strategy. The expected utility of the partitioned coalition A_g in this case, denoted V_p , is given by:

$$V_p(A_g, s) = \max_{P \in M_{A_g}} \left\{ \sum_{A_i \in P} \max\{V_r(A_i, state(A_i, s)), V_t(A_i, state(A_i, s))\} \right\} \quad (5)$$

It is notable that on the right hand side of Equation 5 above we consider only the scenarios where all the members in the sub-coalitions formed resume or terminate the search. This is principally because $V_p(A_i, state(A_i, s))$ has already been taken into account, when we go over all possible partitions (as we choose $\max_{P \in M_{A_g}}$).

The optimal strategy is the mapping $(A_g, s) \rightarrow \{resume, terminate, P\}$, maximizing the expected utility $V^*(A_g, s)$ which can now be formulated as:⁹
 $V^*(A_g, s) = \max\{V_r(A_g, s), V_t(A_g, s), V_p(A_g, s)\}$. This can also be expressed in a more efficient manner as:

$$V^*(A_g, s) = \begin{cases} \max\{V_r(A_g, s), V_t(A_g, s)\} & \text{if } |A_g| = 1 \\ \max\{V_r(A_g, s), \max_{A_i, A_j} \{V_{A_i}^* + V_{A_j}^*\}\} & \text{otherwise} \end{cases} \quad (6)$$

where $V_{A_i}^*$ and $V_{A_j}^*$ are the expected utilities of coalitions A_i and A_j that fully partition coalition A_g , i.e., $V_{A_i}^* = V^*(A_i, state(A_i, s))$, $V_{A_j}^* = V^*(A_j, state(A_j, s))$, $A_i \cup A_j = A_g$, $A_i \cap A_j = \emptyset$, $i \neq j$, $A_j \neq \emptyset$ and $A_i \neq \emptyset$. Notice that in the above equation for calculating $V^*(A_g, s)$ we simplified the calculation to include only size-two partitions (i.e., partition into two sub-coalitions) of coalition A_g . This is because every partition $P \in M_{A_g}$ has already been taken into account recursively as part of the definition of $V^*(A_i, s)$ and $V^*(A_j, s)$. Furthermore, the latter definition used for representing a partition (only size-two partitions) also covers the option in which all agents terminate the search.

While the above analysis relies on coalition costs associated with search, one may attempt to also investigate models where the partitioning itself (into sub-coalitions) is associated with some cost (e.g., the creation of any new coalition has a fixed cost). Such an extension of the model is straightforward and merely requires the subtraction of the applicable partitioning cost elements in Equation 5, leaving the remaining of the analysis unchanged. For example, consider a cost $c_{split}(k)$ induced whenever the coalition partitions into k sub-coalitions.¹⁰ Here, Equation 5 transforms into:

$$V_p(A_g, s) = \max_{P \in M_{A_g}} \left\{ \sum_{A_i \in P} \max\{V_r(A_i, state(A_i, s)), V_t(A_i, state(A_i, s))\} - c_{split}(|P|) \right\} \quad (7)$$

⁹ It is notable that attempting to solve the problem using dynamic programming requires the use of a matrix storing the optimal values of all possible coalitions ($2^{|A_g|}$) over all possible states, not to mention that the solution for each case is not polynomial, thus the complexity is immense.

¹⁰ In this example the cost $c_{split}(k)$ applies to scenarios where at least one sub-coalition continues its search. The number of resulting sub-coalitions in each search round, k , includes also one sub-coalition containing all the agents that terminate the search in this search round (if the search termination strategy applies to at least one agent in the coalition).

4.2 New Challenges and Strategy Characteristics in VSCS

As expected, eliminating the constraint of keeping the coalition's structure fixed, and allowing the coalition to partition itself as part of its strategy space result in several inherent changes in the optimal search strategy's characteristics in comparison to the model of a fixed coalition structure. In particular, an important set of features that are used for efficiently extracting the optimal strategy and overcoming the computational complexity in the FSCS model [14, 38, 8] do not necessarily hold in our model. In the following paragraphs we discuss and illustrate these changes.

Similar to any other general cooperative search model, the incentive to search in the form of coalitions in our model derives from the potential opportunity to reduce the cost of the search associated with obtaining and reviewing opportunities for the searching agents. Nevertheless, in our model, the partition to be used in the optimal strategy is not necessarily the (intuitive) division associated with the minimal search cost per search stage. This non-intuitive result is related to the fact that partitioning decisions are affected both by the coalition's search cost structure and the heterogeneity level between the coalition members' utility functions, as illustrated in the following example.

Example 2 Consider the following simple environment.

Environment 2 Three agents $\{a_1, a_2, a_3\}$ are evaluating the usefulness of a cooperative search in a market associated with 3 types of opportunities $\{\mathbf{o}_1, \mathbf{o}_2, \mathbf{o}_3\}$, equally distributed ($p(\mathbf{o}_i) = 0.33, \forall i$), where each opportunity yields a utility 1 for one of the agents and 0 for the others (e.g., $U_i(\mathbf{o}_j) = 1 \forall i = j, U_i(\mathbf{o}_j) = 0 \forall i \neq j$). The search cost of a single agent is $c(1) = 0.12$, a coalition of 2 agents is associated with a search cost $c(2) = 0.15$ and a cooperative search with a coalition of size 3 is associated with a cost $c(3) = 0.3$.

For this environment, the partition of the agents into sub-coalitions that minimizes the search cost is the one where one of the agents searches by itself, and the two other agents search cooperatively (i.e., $c(1) + c(2) < c(3) < c(1) + c(1) + c(1)$). Nevertheless, a detailed calculation according to Equations 1-6 reveals that the best strategy, yielding the maximum expected utility, is to start the search as a coalition of 3 agents (i.e., with a larger overall search cost), and then partition as necessary according to the optimal strategy. The expected utility of the optimal partition is 2.008, while the expected utility of the partitions that yield the lowest search costs are: 1.885 for $\{a_1\}, \{a_2, a_3\}$, 1.948 for $\{a_2\}, \{a_1, a_3\}$ and 1.9 for $\{a_3\}, \{a_2, a_1\}$. Therefore, when considering potential alternatives in order to extract the optimal search strategy one needs to consider all possible partitioning options and cannot eliminate specific configurations even if another specific partitioning of the coalition that is associated with a smaller search cost can clearly be identified. Formally, even in a scenario where $c(|A_g|) > c(|A_i|) + c(|A_j|)$, $A_i \cup A_j = A_g$, coalition A_g can not be a-priori eliminated from being the optimal coalition with which to continue the search.

An additional complexity in our VSCS model (in comparison to the FSCS which does not include the partitioning option) derives from the absence of the inter-state consistency that characterizes the FSCS model. In the FSCS model, the optimal strategy (in terms of terminating or resuming the search) is fully consistent with former states. Namely, if the coalition's optimal strategy in state s is to resume the search then it is also its strategy in any state s' leading to state s . Contrarily, in VSCS, reaching a potential state s where according to the optimal strategy a specific member of coalition A_g should resume the search, does not indicate that the optimal strategy for this coalition member is to resume the search in all possible former states s' (for whom state s is a potential future state). This characteristic

of the model prevents us from using a simple backward induction mechanism to infer the coalition's strategy as in the FSCS model. Similarly, we cannot infer the coalition's best structure based on the optimal structure found in former stages of a backward induction based mechanism. We present the following example in order to illustrate this inference.

Example 3 Consider the following environment:

Environment 3 Three agents $\{a_1, a_2, a_3\}$ are evaluating the usefulness of a cooperative search in a market associated with 3 types of opportunities $\{\mathbf{o}_1, \mathbf{o}_2, \mathbf{o}_3\}$, equally distributed ($p(\mathbf{o}_i) = 0.33, \forall i$). The search costs are $c(1) = 0.2$, $c(2) = 0.25$ and $c(3) = 0.4$. The agents' utilities are given in the following table:

Opportunities	a_1	a_2	a_3
\mathbf{o}_1	0.1	0.8	0.6
\mathbf{o}_2	0.2	0.82	0.9
\mathbf{o}_3	1	1	1

Now consider the coalition $A_g = \{a_1, a_2, a_3\}$ when in states $s = \{\mathbf{o}_2\}$ and $s' = \{\mathbf{o}_1\}$. Obviously state s' is a possible former state of s , since when the coalition's current state is s' and it encounters opportunity \mathbf{o}_2 its state changes to s (according to the state definition). Thus, knowing the coalition's optimal strategy when in state s is a precondition for finding the coalition's optimal strategy at state s' . A detailed calculation (according to Equations 1-6) reveals that coalition A_g 's optimal strategy when in state s , is having agents a_1 and a_2 resume their search (cooperatively) and agent a_3 terminate its search (the overall expected utility in this case is 2.15). Nevertheless, the fact that agents a_1 and a_2 resume their search in state s , does not necessarily indicate (as in the FSCS model) that this is their optimal strategy when in state s' . For example, in our environment, the maximum expected utility that can be achieved in a configuration by which agents a_1 and a_2 resume their search when in state s' is through coalition $\{a_1, a_2, a_3\}$, yielding 1.975. However, a detailed calculation (according to Equations 1-6) reveals that coalition A_g 's optimal strategy when in state s' , is having agents a_3 and a_1 resume the search cooperatively, whereas agent a_2 should terminate the search (the overall expected utility in this case is 2.075).

Example 3, reflects an additional inherent difference between the optimal strategy structure in the model with the partitioning option (VSCS) and the one which enforces a fixed structure (FSCS), associated with environments where the agents in the searching coalition have correlated preferences. The agents described in Example 3 have correlated preferences (i.e., the agents always produce the same ordinal sorting of the opportunities available in the market). A correlated preferences scenario in the FSCS model yields a stationary optimal search strategy equivalent to the one used by a single agent [28]. According to this strategy the search is constantly resumed until an opportunity associated with an ordinal ranking equal or greater than a preset ranking is reached. Nevertheless, as illustrated in Example 3, the coalition's optimal strategy is not stationary in the VSCS model. Consequently when we have the favorable option of dividing the coalition into sub-coalitions, where agents have correlated preferences, we cannot apply such an analogy in order to simplify computational aspects as in the FSCS model.

Furthermore, even in the case of fully homogeneous agents in the coalition, the use of a single agent equivalent search (as in the FSCS model) is implausible, as illustrated in the following example:

Example 4 Consider the following environment:

Environment 4 Three agents $\{a_1, a_2, a_3\}$ are evaluating the usefulness of a cooperative search in a market associated with 2 types of opportunities $\{\mathbf{o}_1, \mathbf{o}_2\}$, equally distributed ($p(\mathbf{o}_1) = p(\mathbf{o}_2) = 0.5$). The agents' utilities are 0 for opportunity \mathbf{o}_1 and 0.5 for opportunity \mathbf{o}_2 . The search costs are $c(1) = 0.25$, $c(2) = 0.3$ and $c(3) = 0.8$. The coalition's optimal strategy (computed according to Equations 1-6), when starting the search, is to have only two agents resume the search cooperatively and one agent terminate the search. Therefore, despite the fact that the agents are fully homogeneous, we cannot apply the single agent approach.

Overall, in spite of the calculation complexity of the proposed VSCS model, its advantage in terms of the coalition's performance is obvious as the following proposition states.

Proposition 1 *The new cooperative search model (VSCS) is a generalization of both the FSCS and SAS models and weakly dominates them in terms of the overall performance achieved (i.e., may improve the overall expected utility but never worsen it in comparison to the other two methods).*

The proof of the proposition is quite straightforward. FSCS and SAS are both specific cases of the VSCS where the coalition always chooses to resume the search in its original structure or partitions into a set of coalitions of size one (i.e., single agents), respectively. Therefore, if one of these other two search mechanisms produces the maximum utility for a given environment, then the coalition using VSCS will adopt this structure. Notice that between the two methods, FSCS and SAS, none generally dominates the other (but rather the selection of the optimal one is environment-dependent, as demonstrated in [38]). The only advantage of these two methods in comparison to the VSCS is in terms of the computational complexity of extracting the optimal strategy, due to the specific strategy characteristics that do not necessarily hold in the VSCS, as illustrated above. Nevertheless, in many common environments even the extended calculation complexity associated with the VSCS optimal strategy extraction can be overcome and reduced to a level similar to the one obtained for the FSCS model, as we demonstrate in the next section.

5 Reducing Complexity in VSCS

For the general case, the VSCS optimal strategy extraction is exponential in the number of agents. In the remaining of the paper we focus on environments characterized by fixed or non-increasing marginal search costs (formally, described as: $c(n+2) - c(n+1) \leq c(n+1) - c(n), \forall n \geq 0$) to illustrate how this computational complexity can be substantially reduced to become polynomial in the number of agents (Lemma 2). The non-increasing marginal search cost structure is highly common in MAS and can be found in a wide range of environments [26,37]. This is typically due to the fact that most of the coalition overhead is associated with communication. A typical example of this is where in each stage of the search one of the coalition members (or a representative agent) conducts the search on behalf of the coalition. The agent conducting the search needs to send the results to the other agents. The other agents do not have to communicate among themselves, therefore the search cost is at most linear and depends on the number of coalition members. In such environments, as we prove and demonstrate in the following paragraphs, many of the computational complexities induced by allowing coalition partitioning can be overcome. We begin by introducing the following Lemma 1 for environments of this kind, which lays the foundations for our algorithmic-based solution to the problem.

Lemma 1 *In environments characterized by non-increasing marginal search cost, for any state s of a coalition A_g there exists an optimal strategy that involves at most one sub-coalition that resumes the search.*

The above Lemma 1 suggests that any strategy in which the coalition partitions in a way that two or more sub-coalitions (or individual agents) resume their search in the following search round is weakly dominated by a strategy in which k ($k \leq n$) coalition members terminate their search at the current stage while the remaining $n - k$ coalition members continue the search cooperatively. Though the proof below is quite detailed, it is intuitive. If the optimal strategy is to have two sub-coalitions searching in parallel then merging them into one coalition for just one search stage and then having them return to their initial coalition structures will obviously yield a better performance (since the expected utility will remain the same whereas the aggregated cost of such a move can only decrease).

Proof Assume that according to the optimal strategy, coalition A_g partitions into at least two sub-coalitions that resume the search separately. Here, two scenarios may apply: (1) Each coalition, separately, draws the opportunity it encounters from distribution P . (2) There is a list of infinite opportunities $\mathbf{o}^1, \mathbf{o}^2, \dots$ that were taken randomly from distribution P . All coalitions in their i^{th} search stage encounter the same opportunity, \mathbf{o}^i . The probability that a coalition will encounter a given opportunity is exactly the same in both scenarios. Moreover each coalition's utility depends merely on the other opportunities it encounters and doesn't depend on the other coalitions' utility. Therefore each coalition has the same expected utility in both scenarios.

Given this observation, compare the case in which after coalition A_g partitions into sub-coalitions, each sub-coalition resumes the search according to the second scenario with the case that all sub-coalitions of A_g jointly conduct this search. Obviously the sub-coalitions can only benefit from searching together since their search costs can only decrease in the cooperative setting. However after conducting the search in the current search round each sub-coalition can still make the same decisions as in the second scenario. This applies to all cases where a coalition partitions into sub-coalitions that resume the search separately.

5.1 Optimal Cooperative Search Strategy

The immediate implication of Lemma 1 is that in each state the coalition needs to decide on the agents that will continue the search cooperatively and the agents that will terminate the search. Since all agents that resume the search conduct the search together, there is no need for them to decide which coalition to join.

Our analysis suggests a simple mechanism to determine the agents who will cooperatively continue the search at each state for any given coalition. For this purpose we introduce several supporting definitions and notations, applicable for coalition A_g when in state s :

- $V_{one}(s, a_j)$ - the additional expected utility (without incorporating the search cost) that agent $a_j \in A_g$ obtains from terminating its search after conducting one additional search stage rather than terminating the search in the current state s . Thus:

$$V_{one}(s, a_j) = \sum_{U_j(\mathbf{o}) \geq U_j(\mathbf{o}_{a_j}^s)} p(\mathbf{o})(U_j(\mathbf{o}) - U_j(\mathbf{o}_{a_j}^s)) \quad (8)$$

where $\mathbf{o}_{a_j}^s$ is the opportunity, among those known to the coalition in state s , that maximizes agent a_j 's utility.

- $A_{order} = (a'_1, \dots, a'_{|A_g|})$ - the list of agents in A_g sorted in a descending order according to their V_{one} values.
- $A_r = \{a'_1, \dots, a'_k\}$, $A_t = \{a'_{k+1}, \dots, a'_n\}$ - a partition of the sorted list A_{order} , where $k \leq n$ is the first index in A_{order} satisfying both conditions: (C1) $\sum_{j=1}^k V_{one}(s, a'_j) > c(k)$; and (C2) $\nexists i, i > k$ that satisfies

$$\sum_{j=k+1}^i V_{one}(s, a'_j) > c(i) - c(k) \quad (9)$$

If k does not exist then $A_r = \emptyset$ and $A_t = A_g$.

In the above definitions the condition (C1) is used to ensure that the incremental expected utility encapsulated in one additional search stage is greater than the search cost of sub-coalition A_r . The second condition (C2) ensures that the additional utility obtained from moving any subset of A_t to A_r results in a non-positive expected net utility. At this point, we have all the necessary tools to establish the following theorem.

Theorem 1 *The optimal strategy of coalition A_g when in state s is to have the agents in A_r resume the search cooperatively and have the rest of the agents in A_g (i.e., the agents in A_t) terminate the search.*

The general sketch of the proof for Theorem 1 begins by proving that it is sufficient to consider the calculated value V_{one} for determining the optimal partition (rather than use the actual additional expected utility of each agent when using its optimal strategy given the option to resume its search in future states). The proof for this is achieved by showing that each agent's marginal expected utility obtained from resuming the search decreases throughout the search whereas the marginal cost of adding the agent to the coalition that keeps searching can only increase throughout the search (given the search cost structure and the fact that the coalition size throughout the search can only decrease). Therefore if it is not beneficial for the agent to resume its search in the current state given the V_{one} criteria above then this is also the case when using the optimal future strategies. Next we prove that the optimal strategy for all agents in A_r is to resume their search cooperatively as one coalition. This is achieved by showing that under condition (C2), a scenario by which one of the agent's expected additional utility from resuming the search is smaller than its own induced cost will not exist.

Proof We begin by proving that all agents in A_t should terminate the search. We assume by way of contradiction that a subgroup $A_1 \subseteq A_t$, $A_1 \neq \emptyset$ exists such that according to the optimal strategy subgroup A_1 should resume the search in state s . Let subgroup A_2 ($A_2 \subseteq A_1$ and $A_2 \neq \emptyset$), be a coalition of agents and let s' be a state in which: (1) according to the optimal strategy it is possible for A_2 to reach s' during the search. (2) According to the optimal strategy A_2 resumes the search when in s' and terminates the search when in all future states of s' . State s' must exist since the strategy in s is for all agents in A_1 to resume the search and the strategy in a state in which each agent obtains the maximum utility (i.e., when the coalition has found the most appealing opportunities in the market for all its members) is for all agents to terminate the search. Since the number of states is finite we can conclude that such a state exists. Since $k = |A_r|$ satisfies condition C2, $\forall i > k$ (and thus for $i = |A_2| + k = |A_2| + |A_r|$), we obtain:

$$\sum_{j=|A_r|+1}^{|A_2|+|A_r|} V_{one}(s, a'_j) \leq c(|A_2| + |A_r|) - c(|A_r|) \quad (10)$$

Since $a^{|A_r|+1}, \dots, a^{|A_2|+|A_r|}$ are the $|A_2|$ agents with the highest $V_{one}(s, a'_j)$ values in A_r , then the following holds:

$$\sum_{a_j \in A_2} V_{one}(s, a_j) \leq \sum_{j=|A_r|+1}^{|A_2|+|A_r|} V_{one}(s, a'_j) \quad (11)$$

From 10 and 11 we obtain:

$$\sum_{a_j \in A_2} V_{one}(s, a_j) \leq c(|A_2| + |A_r|) - c(|A_r|) \quad (12)$$

Given a state s , a coalition A_g and an arbitrary future state s' that A_g can reach during its search, the following is satisfied: $V_{one}(s', a_j) \leq V_{one}(s, a_j)$, $\forall a_j \in A_g$. This claim stems directly from the calculation of $V_{one}(s, a_j)$ in Equation 8 and from the fact that $U_j(\mathbf{o}'_{a_j}) \leq U_j(\mathbf{o}_{a_j})$, $\forall a_j \in A_g$. From this claim and Equation 12 we find that

$$\sum_{a_j \in A_2} V_{one}(s', a_j) \leq c(|A_2| + |A_r|) - c(|A_r|) \quad (13)$$

Since all agents in A_2 terminate the search in all future states of s' , the left term in Equation 13 is the additional expected utility of a search, if all agents in A_2 resume the search when in s' . The term on the right hand side of Equation 13 is a lower bound to the cost (of a search in state s') when adding all agents in A_2 to the coalition that resumes the search. This is because at state s' at most $|A_r|$ agents will resume the search in addition to the agents in A_2 . This along with the fact that the marginal cost does not increase suggests that the cost of adding the agents in A_2 to the coalition that resumes the search is at least $c(|A_2| + |A_r|) - c(|A_r|)$. Therefore the overall utility of coalition A_g would not increase if all agents in A_2 were to resume the search, which is a contradiction.

Now we only have to prove that according to the optimal strategy all agents in A_r resume the search at state s . We assume by way of contradiction that according to the optimal strategy there is a coalition $A_3 \subseteq A_r$, $A_3 \neq \emptyset$ that terminates the search at state s . We use A_4 to denote the complementary set (i.e., according to the optimal strategy the agents in A_r resume the search when in state s), $A_4 = A_r \setminus A_3$.

Since $|A_r|$ is the first index that satisfies both conditions C1 and C2, $|A_4|$ does not satisfy either C1 or C2. Assume $k = |A_4|$ does not satisfy C1. From this assumption and the fact that $k = |A_r|$ satisfies C1 we obtain:

$$\sum_{j=|A_4|+1}^{|A_r|} V_{one}(s, a_j) > c(|A_r|) - c(|A_4|) \quad (14)$$

Since agents $a^{|A_4|+1} \dots a^{|A_r|}$ are $|A_r| - |A_4| = |A_3|$ agents with the lowest V_{one} values in A_r , the following holds:

$$\sum_{a_j \in A_3} V_{one}(s, a_j) \geq \sum_{j=|A_4|+1}^{|A_r|} V_{one}(s, a'_j) \quad (15)$$

From 14 and 15 we obtain:

$$\sum_{a_j \in A_3} V_{one}(s, a_j) > c(|A_r|) - c(|A_4|) \quad (16)$$

On one hand the right term in Equation 16 gives an upper bound to the cost of a search in state s for adding all agents in A_3 to the coalition that resumes the search. This is because

all agents in A_4 resume the search. Therefore the size of the coalition that resumes the search is at least $|A_4|$. From this and the fact that the marginal cost does not increase, the cost of adding the agents in A_3 to the coalition that resumes the search is at most $c(|A_r|) - c(|A_4|)$. On the other hand the left term in Equation 16 gives a lower bound to the additional expected utility of a search in state s , if the agents in A_3 resume the search (since the expected utility according to the optimal search strategy is at least the expected utility when agents in A_3 conduct a single search). Therefore the expected utility if the agents in A_3 resume the search is higher than the search cost of adding the agents in A_3 to the coalition that resumes the search. Thus the overall expected utility of coalition A_g will increase, if all agents in A_3 resume the search. This is a contradiction to the assumption made.

Now assume $k = |A_4|$ does not satisfy C2. As a result $\exists i > |A_4|$ satisfying:

$$\sum_{j=|A_4|+1}^i V_{one}(s, a'_j) > c(i) - c(|A_4|) \quad (17)$$

Let A_5 be the set of $i - |A_4|$ agents with the highest V_{one} values in $A_r \cup A_3$. Since $A_g = A_4 \cup A_r \cup A_3$, at most $|A_4|$ agents can have a higher utility than $|A_5|$ members, i.e.,

$$\sum_{a_j \in A_5} V_{one}(s, a_j) \geq \sum_{j=|A_4|+1}^i V_{one}(s, a'_j) \quad (18)$$

From 17 and 18 we obtain

$$\sum_{a_j \in A_5} V_{one}(s, a_j) > c(i) - c(|A_4|) \quad (19)$$

Thus for similar considerations as above, we can derive that the overall expected utility of coalition A_g will increase, if all the agents in A_5 resume the search. Again this is a contradiction to the initial assumption made.

Before presenting an algorithm that is based on Theorem 1 for computing the coalition's optimal strategy we illustrate this theorem with the following example:

Example 5 Suppose there are 4 agents $\{a_1, a_2, a_3, a_4\}$ conducting the search in a market associated with 4 types of opportunities $\{\mathbf{o}_1, \mathbf{o}_2, \mathbf{o}_3, \mathbf{o}_4\}$. The search cost associated with a coalition of n agents is given by $c(n) = 0.4ln(n+1)$. The utilities perceived by each agent from each opportunity as well as the distribution of the opportunities are given in the following table:

Opportunities	probability	a_1	a_2	a_3	a_4
\mathbf{o}_1	0.2	5	4	7	8
\mathbf{o}_2	0.1	3	4.4	8.5	7.5
\mathbf{o}_3	0.1	9	5	3	8.5
\mathbf{o}_4	0.6	10	4	6	5

In order, for example, to find the optimal strategy for the coalition $A_g = \{a_1, a_2, a_3, a_4\}$ when the only known opportunity is \mathbf{o}_1 (i.e., $s = \{\mathbf{o}_1\}$) we first need to compute the value $V_{one}(\{\mathbf{o}_1\}, a_j) \forall a_j \in A_g$:

$$V_{one}(\{\mathbf{o}_1\}, a_1) = p(\mathbf{o}_3)(U_1(\mathbf{o}_3) - U_1(\mathbf{o}_1)) + p(\mathbf{o}_4)(U_1(\mathbf{o}_4) - U_1(\mathbf{o}_1)) = 0.1(9 - 5) + 0.6(10 - 5) = 3.4$$

$$V_{one}(\{\mathbf{o}_1\}, a_2) = p(\mathbf{o}_2)(U_2(\mathbf{o}_2) - U_2(\mathbf{o}_1)) + p(\mathbf{o}_3)(U_2(\mathbf{o}_3) - U_2(\mathbf{o}_1)) = 0.1(4.4 - 4) + 0.1(5 - 4) = 0.14$$

$$V_{one}(\{\mathbf{o}_1\}, a_3) = p(\mathbf{o}_3)(U_3(\mathbf{o}_3) - U_3(\mathbf{o}_1)) = 0.1(8.5 - 7) = 0.15$$

$$V_{one}(\{\mathbf{o}_1\}, a_4) = p(\mathbf{o}_2)(U_4(\mathbf{o}_2) - U_4(\mathbf{o}_1)) = 0.1(8.5 - 8) = 0.05$$

Consequently, the sorted agents list is $A_{order} = (a'_1, a'_3, a'_2, a'_4) = (a_1, a_3, a_2, a_4)$.

We start by checking whether the two conditions C1 and C2, are satisfied for $k = 1$ and increase the value of k whenever one of the conditions is not satisfied. The process is summarized in the following table:

	Condition C1	Condition C2
$k = 1$	satisfied: $V_{one}(\{\mathbf{o}_1\}, a'_1) = V_{one}(\{\mathbf{o}_1\}, a_1) = 3.4 > c(1) = 0.277$	not satisfied since Equation 9 holds for $i = 3$ (as $V_{one}(\{\mathbf{o}_1\}, a'_2) + V_{one}(\{\mathbf{o}_1\}, a'_3) = V_{one}(\{\mathbf{o}_1\}, a_3) + V_{one}(\{\mathbf{o}_1\}, a_2) = 0.14 + 0.15 = 0.29 > c(3) - c(1) = 0.277$).
$k = 2$	satisfied: $V_{one}(\{\mathbf{o}_1\}, a'_1) + V_{one}(\{\mathbf{o}_1\}, a'_2) = 3.4 + 0.15 = 3.55 > c(2) = 0.439$	not satisfied since Equation 9 holds for $i = 3$ (as $V_{one}(\{\mathbf{o}_1\}, a'_3) = 0.14 > c(3) - c(2) = 0.115$)
$k = 3$	satisfied: $V_{one}(\{\mathbf{o}_1\}, a'_1) + V_{one}(\{\mathbf{o}_1\}, a'_2) + V_{one}(\{\mathbf{o}_1\}, a'_3) = 3.4 + 0.15 + 0.14 = 3.69 > c(3) = 0.554$.	satisfied
$k = 4$	satisfied	satisfied

Since $k = 3$ is the first index that satisfies both conditions C1 and C2, we conclude based on Theorem 1 that $\{a_1, a_2, a_3\}$ should resume the search and $\{a_4\}$ should terminate the search.

5.2 Computation Algorithm

Based on Theorem 1 we present algorithm 1 for computing the optimal search strategy for any coalition A_g given its current state s . The significance of the algorithm is that it enables

Algorithm 1 Computing the optimal strategy for coalition A_g when reaching state s

Input: $U = \{U_1, \dots, U_{|A_g|}\}$ - coalition members' utility functions; s - the coalition's current state, O - set of potential opportunity types in the market; $p(\mathbf{o})$ - opportunity types' probability function; $c(n)$ - search cost function;

Output: (A_r, A_t) - the sub-coalition that needs to resume the search and the complimentary set of agents that needs to terminate the search at the current stage, respectively.

- 1: Generate the sorted descending set $A_{order} = (a'_1, \dots, a'_{|A_g|})$ by computing $V_{one}(s, a_j) \forall a_j \in A_g$ using Equation 8
 - 2: $size \leftarrow 0; V_{one}^{addCoalition} \leftarrow 0$
 - 3: **for** $index = 1$ to $|A_{order}|$ **do**
 - 4: $V_{one}^{addCoalition} \leftarrow V_{one}^{addCoalition} + V_{one}(s, a'_{index})$
 - 5: **if** $V_{one}^{addCoalition} > c(index) - c(size)$ **then**
 - 6: $size \leftarrow index; V_{one}^{addCoalition} \leftarrow 0$
 - 7: **end if**
 - 8: **end for**
 - 9: return $(\{a'_1, \dots, a'_{size}\}, \{a'_{size+1}, \dots, a'_{|A_g|}\})$
-

us to extract the optimal strategy for a coalition without considering all possible states. As stated in Lemma 2 the computational complexity of the proposed algorithm depends merely on the number of market opportunities and it is polynomial in the market opportunities and agents.

Lemma 2 Algorithm 1 returns the optimal strategy of coalition A_g at state s in a polynomial time of $|O|$ and in $|A_g|$.

Proof Let k be the first index that satisfies both conditions C1 and C2. If such a k does not exist, algorithm 1 returns $A_r = \emptyset$ and $A_t = A_g$, since the condition in step 5 is not satisfied for any value of the variable $index$. If such a k exists we prove that at the end of algorithm 1, the value stored in $size$ is equal to k . In order to prove this we first prove that the condition in step 5 is satisfied when $index = k$. Namely, we prove that $V_{one}^{addCoalition} > c(index) - c(size_{ind=k})$, where $size_{ind=k}$ denotes the current size of the coalition that resumes the search (where $index = k$ and the algorithm performs step 5).

In order to prove that $V_{one}^{addCoalition} > c(index) - c(size_{ind=k})$, where $index = k$ we consider two possible scenarios: (1) The current size of the coalition that resumes the search is 0, i.e., $size_{ind=k} = 0$; and (2) The complementary case, i.e., where $size_{ind=k} > 0$. In the first scenario we know that $V_{one}^{addCoalition} = \sum_{j=1}^k V_{one}(s, a'_j)$ and $c(index) - c(size_{ind=k}) = c(k) - 0 = c(k)$. Since k satisfies condition C1, it follows that the condition in step 5 is satisfied.

In the second scenario ($size_{ind=k} > 0$) we assume by way of contradiction that $V_{one}^{addCoalition} \leq c(index) - c(size_{ind=k})$. The fact that k satisfies condition C2 and $V_{one}^{addCoalition} \leq c(index) - c(size_{ind=k})$, implies that $size_{ind=k}$ satisfies condition C2. Moreover, since the set of agents we add in step 6 to the coalition that resumes the search always satisfies the condition in step 5, $size_{ind=k}$ also satisfies condition C1. In conclusion, we have shown that $size_{ind=k}$ satisfies both conditions C1 and C2 which contradicts the fact that k is the first index which satisfies both conditions C1 and C2. Consequently, in both cases $V_{one}^{addCoalition} > c(index) - c(size_{ind=k})$ holds. Therefore in step 6 the algorithm increases the size of the coalition that resumes the search to k . Furthermore, since k satisfies condition C2, the condition $V_{one}^{addCoalition} > c(index) - c(size_{ind=k})$ given in step 5 is not satisfied for $index > k$. Therefore the algorithm does not continue to increase the size of the coalition that resumes the search and its size remains k .

Now we only need to prove that algorithm 1 is polynomial in $|O|$ and in $|A_g|$. The step which affects the time complexity of the algorithm the most is step 1. This step requires computing $V_{one}(s, a'_j)$ for all agents. The complexity of computing $V_{one}(s, a'_j)$ for a single agent requires considering all opportunities in O and its complexity is the order of $|O|$. Consequently the complexity of computing $V_{one}(s, a'_j)$ for all agents is the order of $|O||A_g|$. Finally the complexity of sorting the agents in A_g in order to generate A_{order} is the order of $|A_g| \log(|A_g|)$. The loop in step 3 is performed $|A_g|$ times, where each iteration's complexity (the complexity of steps 4-6) is constant. Consequently, the global complexity of algorithm 1 is the order of $|O||A_g| + |A_g| \log(|A_g|)$.

5.3 Value Calculation

Algorithm 1 enables the extraction of the coalition's optimal partitioning-enabled cooperative search strategy (VSCS). However, alongside knowing how to conduct the search optimally, the agents might also be interested in having efficient computational means for calculating the performance (i.e., the expected overall utility) of the coalition, $V^*(A_g, s)$, when searching cooperatively using the optimal VSCS strategy. In particular, having the ability to calculate $V^*(A_g, s)$ enables and facilitates future payoff division and coalition stability mechanism design related research in the cooperative search domain. Obviously in this case the calculation complexity increases in comparison to the case where only the optimal strategy is needed. Nonetheless, we manage to supply this functionality with the complexity

(similar to the FSCS model) solely dependent on the number of states in S_{A_g} rather than on the number of possible sub-coalitions.

Given state s and coalition A_g we use $resume(A_g, s)$ ($terminate(A_g, s)$) to denote the agents that resume (terminate) the search according to algorithm 1, respectively. The value $V^*(A_g, s)$ can be computed using:

$$V^*(A_g, s) = V_r(resume(A_g, s), s_{res}) + V_t(terminate(A_g, s), s_{ter}) \quad (20)$$

where $s_{res} = state(resume(A_g, s), s)$ and $s_{ter} = state(terminate(A_g, s), s)$. We use algorithm 2 to calculate $V_r(resume(A_g, s), s_{res})$ in the above Equation 20. This algorithm makes use of the following equation which is obtained by replacing $V^*(A_g, state(A_g, s \cup \{\mathbf{o}\}))$ in Equation 4 with the right hand side of Equation 20:

$$V_r(A_g, s) = \frac{\sum_{\mathbf{o} \in O_{improve}^s} p(\mathbf{o})(V_r(A'_r, s'_{res}) + V_t(A'_t, s'_{ter})) - c(|A_g|)}{\sum_{\mathbf{o} \in O_{improve}^s} p(\mathbf{o})} \quad (21)$$

where $s' = state(A_g, s \cup \{\mathbf{o}\})$, $A'_r = resume(A_g, s')$, $A'_t = terminate(A_g, s')$, $s'_{res} = state(A'_r, s')$ and $s'_{ter} = state(A'_t, s')$.

Algorithm 2 *ComputeVResum*(A_g, s) - computing the value of resuming the search when in a given state s , $V_r(A_g, s)$

Input: s - current state; A_g - coalition resuming search; $U = \{U_1, \dots, U_{|A_g|}\}$ - coalition members' utility functions; O - set of potential opportunity types in the market; $p(\mathbf{o})$ - opportunity types' probability function; $c(n)$ - search cost function;

Output: V_{resume} - the collection of stored $V_r(A_g, s)$ values

- 1: $O_{improve}^s \leftarrow$ all opportunities that change A_g 's current state
 - 2: **for all** $\mathbf{o} \in O_{improve}^s$ **do**
 - 3: $s' \leftarrow state(A_g, s \cup \{\mathbf{o}\})$, $A'_r \leftarrow resume(A_g, s')$
 - 4: $A'_t \leftarrow terminate(A_g, s')$, $s'_{res} \leftarrow state(resume(A_g, s'), s')$
 - 5: **if** $V_r(A'_r, s'_{res}) \notin V_{resume}$ **then**
 - 6: $ComputeVResum(A'_r, s'_{res})$
 - 7: **end if**
 - 8: **end for**
 - 9: Compute $V_r(A_g, s)$ using Equation 21
 - 10: Store $V_r(A_g, s)$ in its corresponding place in V_{resume}
 - 11: return V_{resume}
-

The computation of $V_r(resume(A_g, s'), s'_{res})$ is achieved by calling algorithm 2 recursively, where each value calculated by the algorithm is stored in the collection V_{resume} . This prevents redundant calculations where the same values are needed when considering different sequences of opportunities (see step 5 in the algorithm).

Given a coalition A_g and a state s , we prove that the number of coalitions and states whose values are calculated by algorithm 2 (and stored in V_{resume}) can be bounded by $|S_{A_g}|$.

Lemma 3 *Given coalition A_g and state s , the number of coalitions and states calculated by algorithm 2 and stored in V_{resume} during the calculation of $V_r(A_g, s)$ can be bounded by $|S_{A_g}|$.*

Proof Let $|S_{A_g}|$ be the collection of all possible states for coalition A_g . For each $s' \in S_{A_g}$ we denote the sub coalition of coalition A_g that resumes the search in s' by $A'_r = resum(A_g, s')$

and the state of the coalition that resumes the search by $s_r^{s'} = \text{state}(A_g', s')$. We show that for each coalition A_g' and state s' for which the $V_r(A_g', s')$ value was stored during the algorithm execution, there exists a state $s_1 \in S_{A_g}$ such that (1) $A_g' = A_r^{s_1}$; and (2) $s' = s_r^{s_1}$. Therefore the size of the collection V_{resume} is bounded by $|S_{A_g}|$. Given coalition A_g' and state s' for which the algorithm calculates a value in step 6, let $O_{\text{stay}}^{s'}$ be the set of opportunities that does not change coalition A_g' 's current state. The opportunities in $O_{\text{stay}}^{s'}$ are the only opportunities that the coalition could have encountered until the current search stage, otherwise the current state would not be s' . In order to create s_1 we extend s' in a way that s_1 also includes the opportunities that maximize the rest of the agents' ($A_g \setminus A_g'$) utilities when using the opportunities in $O_{\text{stay}}^{s'}$, $s_1 = s' \cup \text{state}(A_g \setminus A_g', O_{\text{stay}}^{s'})$. In order to prove that s_1 satisfies condition (1), we first prove that $A_g' \subseteq \text{resume}(s_1, A_g)$. We assume by way of contradiction that $\exists A_1 \subseteq A_g'$, $A_1 \neq \emptyset$ such that its strategy in s_1 is to terminate the search. We use A_2 to denote the agents in A_g' that resume the search in s_1 , i.e., $A_2 = A_g' \setminus A_1$. Since A_g' 's strategy in state s' according to algorithm 1 is to have all its agents resume the search (i.e., $A_g' = \text{resume}(A_g', s')$), we attain from Theorem 1 that:

$$\sum_{a_j \in A_1} V_{\text{one}}(a_j, s') > c(|A_g'|) - c(|A_2|) \quad (22)$$

From the definition of s_1 it follows that $U_j(\mathbf{o}_{a_j}^{s'}) = U_j(\mathbf{o}_{a_j}^{s_1})$, $\forall a_j \in A_g'$. Therefore:

$$\sum_{a_j \in A_1} V_{\text{one}}(a_j, s') = \sum_{a_j \in A_1} V_{\text{one}}(a_j, s_1) \quad (23)$$

From Equations 22 and 23 we obtain:

$$\sum_{a_j \in A_1} V_{\text{one}}(a_j, s_1) > c(|A_g'|) - c(|A_2|) \quad (24)$$

The term on the left hand side of Equation 24 provides a lower bound for the additional expected utility of continuing the search when in state s_1 , if all the agents in A_1 resume the search in this state. The term on the right hand side of Equation 24 provides an upper bound for the search cost when adding the agents in A_1 to the coalition that resumes the search. This is because the size of the coalition that resumes the search in state s_1 is at least $|A_2| = |A_g'| - |A_1|$. Therefore the additional utility that can be obtained if A_1 members resume the search is greater than the additional search cost. As a result, the coalition's overall utility will increase, if A_1 resumes the search in s_1 , which is a contradiction to the initial assumption.

Next, we prove that $\text{resume}(A_g, s_1) \subseteq A_g'$. We assume by way of contradiction that $\exists A_3 \subseteq A_g \setminus A_g'$, where a sub-coalition $A_3 \neq \emptyset$ resumes the search in state s_1 and all other agents in A_g (that are not in A_3 or in A_g') terminate the search. The fact that algorithm 2 calculates the A_g' value in state s' in step 6, implies that all other agents, including A_3 's members, terminate their search earlier (as the algorithm calculates the value of the coalition that resumes the search). Therefore a state s_2 (a former state of s') exists such that all members of A_3 reach s_2 , whereas a subset of A_3 members, denoted A_4 , terminate the search in state s_2 . Since all agents in A_3 resume the search in state s_1 , all agents in A_4 also resume the search. From Theorem 1 we obtain that

$$\sum_{a_j \in A_4} V_{\text{one}}(s_1, a_j) > c(|A_g'| + |A_3|) - c(|A_3 \setminus A_4| + |A_g'|) \quad (25)$$

State s_1 includes the opportunities that maximize A_4 members' utilities for all the opportunities the coalition A_g could encounter until coalition A_g' reaches state s' including the

opportunities in s_2 . Thus $U_j(\mathbf{o}_{a_j}^{s_1}) \geq U_j(\mathbf{o}_{a_j}^{s_2}), \forall a_j \in A_4$. This, along with the way V_{one} was defined, guarantee the following:

$$\sum_{a_j \in A_4} V_{one}(s_2, a_j) \geq \sum_{a_j \in A_4} V_{one}(s_1, a_j) \quad (26)$$

From Equations 25 and 26 we find that

$$\sum_{a_j \in A_4} V_{one}(s_2, a_j) > c(|A'_g| + |A_3|) - c(|A_3 \setminus A_4| + |A'_g|) \quad (27)$$

The right term of Equation 27 is an upper bound for the cost of a search when in state s_2 if adding the agents from A_4 to the coalition that resumes the search when in state s_2 . This is because the size of the coalition that resumes the search in s_2 is at least the coalition size when in s_1 (there are at least $|A'_g| + |A_3 \setminus A_4|$ that resume the search in state s_2) and the marginal costs do not increase. The left term is a lower bound to the additional expected utility of coalition A_4 for a search in state s_2 . Therefore the overall utility will increase if A_4 resumes the search in s_2 , which is a contradiction. As a result s_1 satisfies condition (1). The fact that state s_1 satisfies condition (2) follows directly from condition (1) and the way s_1 is constructed.

As stated in Lemma 3 the usefulness of algorithm 2 is in the ability to extract a coalition's optimal strategy, given the partitioning capability, without considering all possible coalitions (which is known to be exponential). Moreover the fact that we can determine the coalition strategy without computing the coalition value, can significantly reduce the complexity of algorithm 2. This is because there are future states that we may not reach during the search. For example if the coalition's optimal strategy in a given state is to have all agents terminate the search then the computation of $V^*(A_g, s)$ is immediate (equals to $V_t(A_g, s)$) and does not depend on any future state of s .

Finally we note that the superiority of the VSCS model over the FSCS model (in environments with marginal cooperative search costs) is when the agents are heterogenous (even when agents have correlated preferences). However, as the following lemma states, when the agents are fully homogenous the coalition's optimal strategy returned by the VSCS model is equal to the coalition's optimal strategy returned by the FSCS model.

Lemma 4 *Given state s and coalition A_g , if all agents are fully homogeneous then both VSCS and FSCS models return the same optimal strategy.*

Proof Since all agents are homogeneous, they share the same $V_{one}(s, a_j)$ values for each state s . We consider the two following cases: (1) The agent's marginal utility for conducting a single search exceeds the agent's average share of the coalition's search cost, i.e., $V_{one}(s, a_j) > c(|A_g|)/|A_g|$; and (2) The complementary case, i.e., $V_{one}(s, a_j) \leq c(|A_g|)/|A_g|$. In the former case, condition C2 is not satisfied for each $ind = 1, \dots, |A_g|$ (since the marginal costs are non-increasing). Therefore according to Theorem 1 the optimal strategy is to have all agents resume the search at state s . Moreover, the overall additional utility of the coalition's members from conducting a single search exceeds the coalition's search cost since $V_{one}(s, A_g) = \sum_{j=1}^{|A_g|} V_{one}(s, a_j) > \sum_{i=1}^{|A_g|} c(|A_g|)/|A_g| = c(|A_g|)$. Therefore the optimal strategy according to the FSCS model is to have all agents resume the search at state s (the same search strategy as in VSCS model).

In the latter case where $(V_{one}(a_j, s) \leq c(|A_g|)/|A_g|)$, no index satisfies condition C1. Therefore all agents terminate the search at state s in the VSCS model. Moreover the overall additional utility of the coalition's members from conducting a single search does not exceed

the coalition's search cost, $V_{one}(s, A_g) = \sum_{j=1}^{|A_g|} V_{one}(s, a_j) \leq \sum_{i=1}^{|A_g|} c(|A_g|)/|A_g| = c(|A_g|)$. As for the FSCS, a proof that the optimal coalition's search strategy in such scenarios when reaching state s is to have all agents terminate the search is given in [26].

6 Illustrative Comparison

In this section we exemplify certain properties of the VSCS model and illustrate the differences in the optimal strategy structure and in the overall performance between the proposed VSCS model and the FSCS and SAS models. Notice that given Proposition 1 the superiority of the VSCS over the FSCS and the SAS models is unquestionable. Therefore our goal is mainly to demonstrate different aspects of the optimal search strategies in each of the different search methods. Finally after demonstrating these aspects we illustrate the advantages of the proposed model over currently known cooperative and individual search techniques, using an environment based on authentic settings. These settings contain thousands of opportunity types. Nonetheless, with the use of the algorithms that were introduced in the previous section, we can easily extract the coalition's optimal strategy, demonstrating the applicability of the proposed search method.

6.1 Demonstration of the Search Process in the VSCS model

We begin by demonstrating the general search process that takes place in the VSCS model using the following synthetic environment.

Environment 5 A coalition of 5 agents $A_g = \{a_1, a_2, a_3, a_4, a_5\}$, searching for opportunities in a marketplace characterized by three types of opportunities uniformly distributed. The agents' utility for the different opportunities as well as the coalition's search cost as a function of its size are given in table below:

The search cost of a single agent is $c(1) = 0.22$, a coalition of 2 agents is associated with a search cost $c(2) = 0.25$, a coalition of 3 agents is associated with a search cost $c(3) = 0.4$, a coalition of 4 agents is associated with a search cost $c(4) = 0.55$ and a coalition of 5 agents is associated with a search cost $c(5) = 0.56$.

Opportunities	a_1	a_2	a_3	a_4	a_5
\mathbf{o}_1	1	1	2	2	2
\mathbf{o}_2	0.5	0.5	0.5	0.5	2
\mathbf{o}_3	2	2	1	1	2
Coalition size	1	2	3	4	5
Search Cost (\$)	0.22	0.25	0.4	0.55	0.56

Note that the marginal cooperative search cost in Environment 5 has a mixed structure, since on the one hand $c(3) - c(2) > c(2) - c(1)$ and on the other $c(5) - c(4) < c(4) - c(3)$. Therefore, we expect the coalition to possibly partition itself to more than one active sub-coalition throughout cooperative search.

Figure 1 shows all potential coalition structures and states that can occur during the search process that takes place in Environment 5 (when the coalition uses its optimal strategy). Each active coalition and its current state is represented by an oval. The squares detail the coalition members that terminate the search and the opportunities that each of them exploits. For example, the optimal strategy of coalition $\{a_1, a_2, a_3, a_4, a_5\}$ at the beginning of its search, where it encounters opportunity \mathbf{o}_1 is to have agents a_3, a_4 and a_5 terminate the

search (as shown in the square) and agents $\{a_1, a_2\}$ resume the search cooperatively, where their current state is \mathbf{o}_1 (as shown in the oval). To simplify the figure we do not include transitions (i.e., the finding of new opportunities) that do not change the coalition's state. As can be noticed from Figure 1, the coalition exploits the ability to partition itself in order to improve its expected performance. For example agents a_1, a_2 and a_5 terminate their search after they encounter opportunity \mathbf{o}_3 , which maximizes their utility. This reduces the coalition size and yields a smaller coalition search cost for the remaining search. Moreover the coalition can reduce its search cost by partitioning itself into sub-coalitions which conduct the search separately. The partitioning can occur in each stage of the search and not necessarily at the beginning of the search (before the coalition starts accumulating opportunities). Indeed as depicted in Figure 1 the optimal strategy of coalition $\{a_1, a_2, a_3, a_4, a_5\}$ is to partition into two sub coalitions $\{a_1, a_2\}$ and $\{a_3, a_4\}$ after it encounters \mathbf{o}_2 (when having no opportunities to begin with).

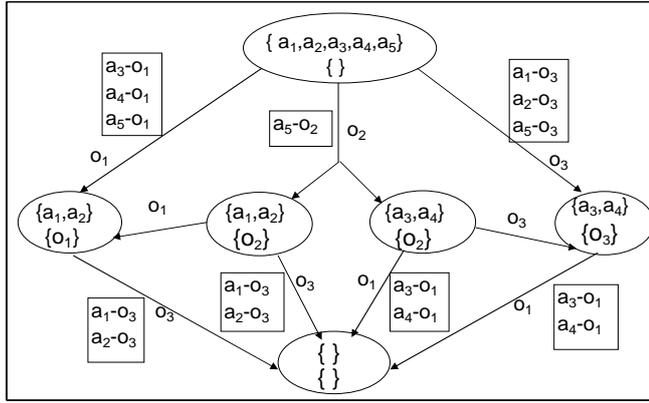


Fig. 1 The coalition's optimal structure and its current state during the search process, when the marginal cooperative search cost has a mixed structure.

Next we demonstrate the search process, when the marginal cooperative search cost has a non-increasing structure (as used for the analysis given in section 5).

Environment 6 A coalition of 4 agents $A_g = \{a_1, a_2, a_3, a_4\}$, searching for opportunities in a marketplace characterized by three types of opportunities uniformly distributed. The search cost associated with n agents conducting the search is $c(n) = 0.4 * \ln(n + 1)$. The agents' utility from the different opportunities are given in the following table:

Opportunities	a_1	a_2	a_3	a_4
\mathbf{o}_1	1	10	4	3
\mathbf{o}_2	3	9.3	12	1
\mathbf{o}_3	10	9.4	2	9

Similar to Figure 1, Figure 2 shows the possible coalition structures that can occur during the search process. Here (according to Lemma 1), in each stage of the search only one sub coalition (at most) resumes the search. Moreover, as this figure illustrates, when the coalition members terminate their search, they do not necessarily exploit the same opportunities. When the sub-coalition $\{a_2, a_3\}$ (which is the reduced coalition after the grand coalition encounters opportunity \mathbf{o}_3) encounters opportunity \mathbf{o}_2 it terminates its search while

its members exploit different opportunities (agent a_2 exploits opportunity \mathbf{o}_3 and agent a_3 exploits opportunity \mathbf{o}_2).

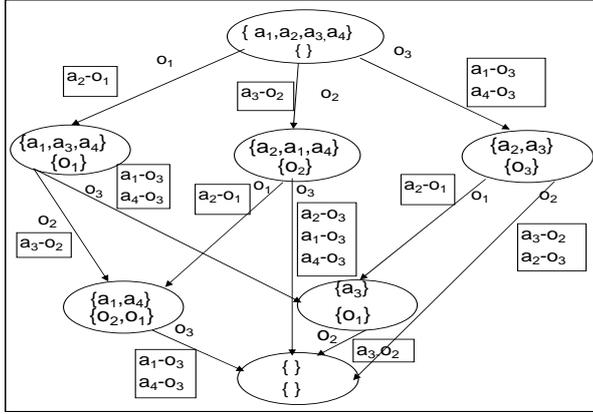


Fig. 2 The coalition's optimal structure and its current state during the search process in Environment 6 (where the marginal cooperative search cost has a non-increasing structure).

Finally we demonstrate the influence of the coalition's partitioning capability in the VSCS model on the coalition's expected overall utility. In order to do so we consider the following environment:

Environment 7 A coalition of 3 agents $A_g = \{a_1, a_2, a_3\}$, conducting the search in a market associated with 3 types of opportunities. The search cost is given by $c(n) = 0.2 + 0.02n^3$. The agents' utility from the different opportunities are given in the following table:

Opportunities	probability	a_1	a_2	a_3
\mathbf{o}_1	0.989	1	2	8
\mathbf{o}_2	0.001	270	270	1
\mathbf{o}_3	0.01	8	1	30

Figure 3 depicts the expected coalition's overall utility with respect to the coalition's partition at the beginning of the search (i.e., the first search stage, before the agent knows about any of the opportunities), assuming all the sub-coalitions created use their optimal strategy. Here, we can see that the coalition's expected utility is also influenced by the agents constituting the sub-coalitions and not only by the size of the sub-coalitions. For example, the three partitions $\{\{a_1\}, \{a_2, a_3\}\}$, $\{\{a_2\}, \{a_1, a_3\}\}$ and $\{\{a_3\}, \{a_1, a_2\}\}$ of the coalition divide into two sub coalitions, one with 2 agents and the other with 1 agent yielding different expected overall utilities. From this figure, we learn that the partition of the coalitions that yields the maximum expected utility is $\{\{a_1, a_2\}, \{a_3\}\}$.

6.2 Influence of Environment Parameters

An important parameter that affects the performance (in terms of the expected utility) of the searching coalition in the VSCS model (as well as in the FSCS model) is the marginal search cost associated with each increase in the coalition's size. In order to demonstrate this affect we use the following environment:

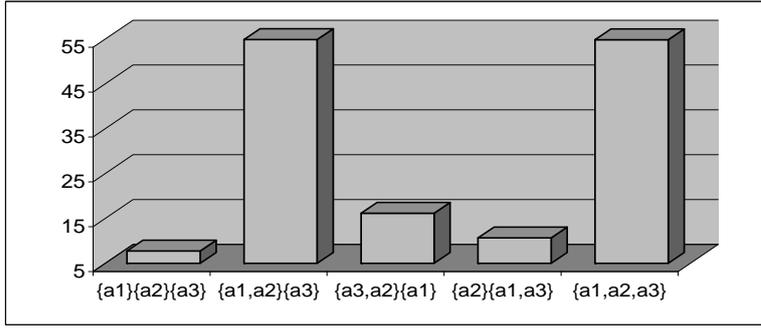


Fig. 3 The overall utility as a function of the initial partition at the beginning of the search.

Environment 8 A coalition of five agents, $A_g = \{a_1, \dots, a_5\}$, searching for opportunities in a marketplace characterized by ten types of opportunities $\mathbf{o}_1, \dots, \mathbf{o}_{10}$ uniformly distributed. The five agents are heterogeneous with respect to the way they evaluate each opportunity type and their utility functions can be described using the function:

$$U_j(\mathbf{o}_i) = \begin{cases} 7 & i = j, j+5 \\ 1 & \text{otherwise.} \end{cases} \quad (28)$$

The search cost of the coalition is assumed to be equal to the sum of some interaction cost (the cost of locating a seller and communicating with her to learn her offer) and communication cost (for transferring the search results to the other $n-1$ agents), i.e., $c(n) = c_{\text{interaction}} + c_{\text{communication}} * (n - 1)$, where $c_{\text{interaction}} = 1$ (i.e., $c_{\text{communication}}$ is the marginal cost when increasing the coalition size by one).

Figure 4 depicts the overall expected utility when using the three search methods (VSCS, FSCS and SAS) as a function of the communication cost $c_{\text{communication}}$ (notice that the agents' performance is not affected by this value in SAS, as it does not require any communication between the agents). As expected, VSCS dominates both methods for any $c_{\text{communication}}$ value. From the graph we observe that the higher the value of $c_{\text{communication}}$ the greater the relative improvement achieved when using the VSCS model in comparison to FSCS. This observation can be explained by the additional reduction achieved in the search cost as a result of resuming the search in a reduced coalition structure instead of keeping the original coalition's size. Moreover in the extreme case, where there is no overhead associated with adding additional members to the coalition (i.e., $c_{\text{communication}} = 0$), the VSCS performance converges to the one exhibited by the FSCS. This is due to the fact that in this case the coalition does not benefit from reducing the search cost as a result of reducing the coalition's size (since adding additional agents does not change the coalition's search cost). Consequently the coalition's optimal search strategy is to keep its original structure. On the other hand as $c_{\text{communication}}$ decreases, a higher improvement is achieved by using the VSCS model in comparison to the SAS model. In the extreme case where the communication cost of adding an additional agent to the coalition that resumes the search is equal to the search cost of an agent searching individually, $c_{\text{communication}} = c_{\text{interaction}}$, the VSCS converges to the SAS model's performance.

As discussed in the previous sections, the ability to partition the coalition results in extended searches. Figure 5 emphasizes the differences in the search extent between the FSCS

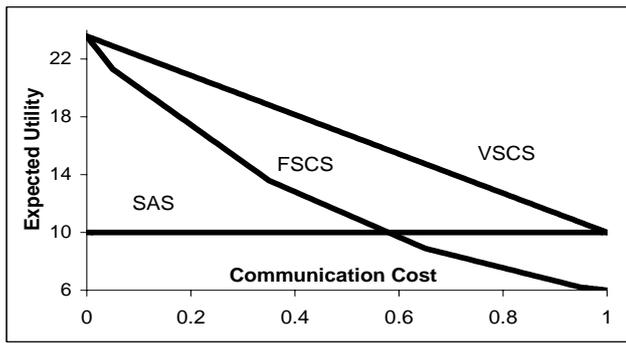


Fig. 4 Overall expected utility for the different models in Environment 8

and the VSCS methods. This figure makes use of Environment 8. In this case, we simulated the search of the initial coalition and kept track of the number of search rounds executed in each method ($c_{communication} = 0.5$). Each column in the graph represents the percentage of coalitions (out of 1000 simulation runs) that lasted the specific search round (measured on the horizontal axis) in the VSCS where the internal division maps the coalition size at that search step as a percentage of the total coalitions. The line curve depicts the appropriate percentage of coalitions continuing the search at that search stage in FSCS. As expected, the adaptive nature of the VSCS allows the coalition to maintain the search far beyond the point where the FSCS coalitions find the option of resuming the search to be non-beneficial and thus terminate their search. Consequently, in VSCS the overall number of opportunities encountered is greater than in FSCS.

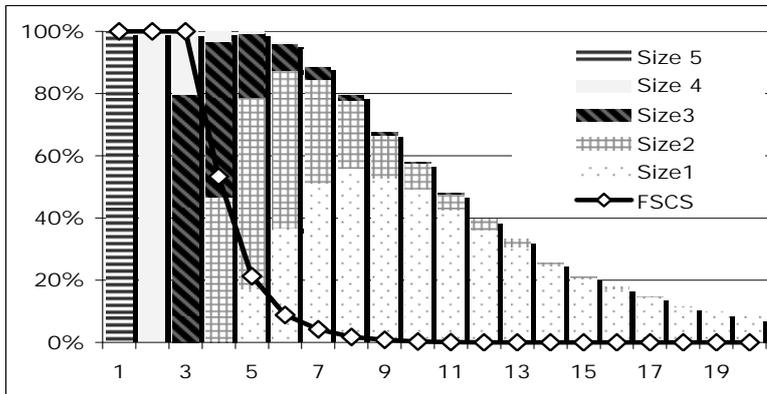


Fig. 5 Extent of search according to the different methods

Similar to the FSCS model, the level of heterogeneity in the utility functions of the different coalition members highly affects the coalition's performance (in terms of the expected utility achieved). In order to demonstrate this we use the following environment.

Environment 9 A coalition of four agents, $A_g = \{a_1, a_2, \dots, a_4\}$, searching for opportunities defined by four attributes, B_1, B_2, B_3 and B_4 , where each attribute has only two possible values (1,2) with an equal probability. The utility functions associated with agents $a_1, a_2,$

a_3 and a_4 are:

$$U_1(\mathbf{o}) = 4(\alpha * B_1 + \frac{1-\alpha}{3} * B_2 + \frac{1-\alpha}{3} * B_3 + \frac{1-\alpha}{3} * B_4)$$

$$U_2(\mathbf{o}) = 4 * (\frac{1-\alpha}{3} * B_1 + \alpha * B_2 + \frac{1-\alpha}{3} * B_3 + \frac{1-\alpha}{3} * B_4)$$

$$U_3(\mathbf{o}) = 4 * (\frac{1-\alpha}{3} * B_1 + \frac{1-\alpha}{3} * B_2 + \alpha * B_3 + \frac{1-\alpha}{3} * B_4)$$

$$U_4(\mathbf{o}) = 4 * (\frac{1-\alpha}{3} * B_1 + \frac{1-\alpha}{3} * B_2 + \frac{1-\alpha}{3} * B_3 + \alpha * B_4)$$

Thus the deviation of the parameter α from the value 0.25 (in which the agents are fully homogeneous) indicates the level of the agents' heterogeneity. The search cost of the coalition is equal to the sum of the interaction cost and the communication cost of reporting the search results to the other $n - 1$ agents as before, i.e., $c(n) = c_{interaction} + c_{communication} * (n - 1)$, where $c_{interaction} = 0.75$ and $c_{communication} = 0.3$.

Figure 6 depicts the overall expected utility in each of the search methods as a parameter of α (which relates the similarity between the utility functions of the agents constituting the coalition). As illustrated in the figure, the greater the heterogeneity level (the larger the deviation of the parameter α from the value 0.25) the higher the improvement achieved when using VSCS in comparison to FSCS. Note that for $\alpha = 0.25$, where the agents are fully homogenous, the overall expected utility of FSCS and VSCS models are equal. This result is consistent with Lemma 4.

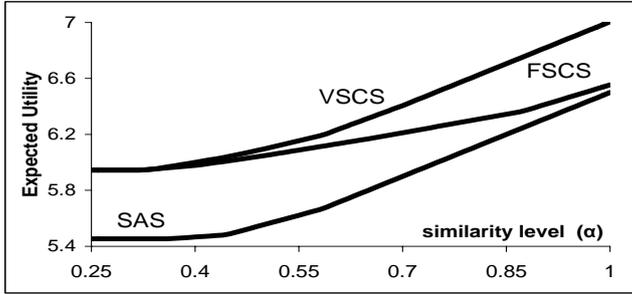


Fig. 6 Overall expected utility as a parameter of the similarity level between the agents' utility functions

6.3 The VSCS Model in an Authentic Environment

While the previous illustrations used synthetic environments for emphasizing various aspects of the proposed model, we wish to illustrate the use of the proposed VSCS model in authentic environment. For this purpose, we formed the following Environment 10 based on opportunities collected over the internet and utility functions that were defined by human searchers we interviewed.

Environment 10 *The searching coalition consists of seven agents interested in buying a calculator. Each agent is associated with a different utility function, based on typical attributes of calculators (price, handled/non-handled, display type, scientific functions, warranty, calculator's company, 2\1 line display, etc.). The utility functions of the different agents were constructed according to real preferences of 7 people (evaluating the different attributes using monetary units) that we interviewed. For example, the worth of a true fraction display was set at \$30 by one of the subjects that was searching for a handled scientific calculator. Similarly, a calculator that does not have these features had no value at all to this individual. The opportunities to buy the calculator in this environment are drawn from a distribution that is based on current offerings over the internet on US-based ecommerce web-sites. Cooperative search is executed by having one of the coalition members conduct*

the search at each stage of the search and informing the search results to the other coalition members. Therefore the search cost of the coalition is equal to the sum of the interaction cost and the communication cost of reporting the search results to the other $n-1$ agents, i.e., $c(n) = c_{interaction} + c_{communication} * (n - 1)$. The interaction cost we used was \$0.01.

Figure 7 depicts the average overall utility over 10000 searches using the three methods as a function of the communication cost (recall that the agents' performance is not affected by this value in the SAS model). As expected, the FSCS performs better than the SAS for some $c_{communication}$ values (and the SAS performs better for others) while the VSCS dominates both methods for any $c_{communication}$ value.

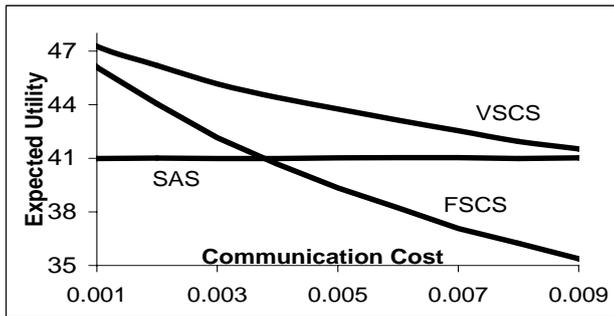


Fig. 7 Overall expected utility for the different models

Even though the above example is based on 3772 opportunity types, it took the coalition only 2 seconds to compute the optimal search strategy (when running on a standard pentium PC, and the algorithm programmed in C++). A typical simulation of a search process draws on average 534 opportunities and lasts less than 2 seconds.

Assuming that the agents are searching cooperatively (or can form such cooperation), the example given illustrates that a better overall expected utility can be obtained using VSCS and that the extraction of the optimal strategy is relatively fast using the algorithms supplied in this paper. Obviously, better examples could have been constructed in order to motivate the formation of the cooperation in the first place (i.e., with complex products and higher variation in utilities). Yet, for the purposes of this section, the use of the calculator example illustrates the different aspects of using VSCS without getting into complex utility elicitation processes (as in the case of more complex real-life environments).

7 Discussion and Conclusions

As shown in the analysis section, the proposed cooperative search model is a generalization of the fixed-structure cooperative search (FSCS) and individual search (SAS) and its use is always favorable when overall expected utility maximization is concerned. Therefore, the option to partition the search, whenever searching cooperatively or having multiple goals for the search, should always be considered. The integration of the new method in search processes is straightforward for the single searcher, when having multiple search goals and easy to implement in the multi-agent case since the ability to maintain an adaptive coalition that can restructure itself is inherent in multi-agent domains [33]. The adaption of the new search method results with a new search strategy, different in its structure in comparison to the optimal strategy used in the fixed structure cooperative search and inherently

different from the strategy used when each agent searches individually. Although we use the electronic marketplace as a framework in this paper, the suggested analysis is general and can be applied to various domains where agents can benefit from cooperative search (e.g., searching a large database of potential candidates to fill several positions). It is notable that the effectiveness of the new method, as any cooperative search, is limited in the case where the things being sought are rare (e.g., a collectible), as in this case conflict of interests, while searching, may arise. Nonetheless, in various settings (e.g., B2C) all search goals can be satisfied by any opportunities found.

The novelty of the analysis given is that it manages to significantly reduce the computational complexity associated with extracting the optimal coalition strategy in environments with a non-increasing marginal search cost. This search cost structure is highly common in MAS environments. The algorithms we supply do not induce any computational increase in comparison to the models without the proposed enhancement. This further increases the potential for actual implementation of cooperative search, especially in the eCommerce domain, where buyer-groups are commonly formed (e.g., for obtaining a volume discount). In this sense, the new method is yet another incentive for cooperation in such markets.

Alongside the contribution to economic search theory (in the context of multi-goal search), the results highly apply to cooperation in multi-agent domains. While the focus of this paper is on finding the coalition strategy that will maximize the overall utility, there are many other aspects of the new model that should be addressed. These include coalition stability, payoff division mechanisms (and in particular when partitioning the coalition along its search) and truth telling [41,43]. Though these were not included within the framework of the current paper, we wish to emphasize that the optimal cooperative search strategy is not influenced by these factors but rather influences them. The analysis of these important issues is based on the ability to properly derive the coalition's utility given its initial specific self structure (i.e., the number of agents it represents and their reported, not necessarily true, utility functions) and the environment within which it operates. By supplying this functionality, we enable extensive important future research and support the integration of various relevant ideas from the rich literature in the area of game theory and MAS research [41,23,34,45] to the proposed model.

Furthermore, while the paper considers the proposed mechanism as a self-contained efficient cooperative search method, we believe there is great potential in applying it as an infrastructure for evaluating heuristics aiming towards further enhancing the cooperative search performance. Such heuristics can consider ad-hoc or occasional communication between formerly partitioned sub-coalitions (and possibly the union of formerly partitioned coalition). These kinds of models are extremely complicated and require heuristic-based approaches. The advantage of the proposed method (VSCS) in this context is threefold. First, it can be used as a baseline for evaluating the performance of these heuristics. Second, its performance can be used as a lower bound for the expected utility encapsulated in any sub-coalition partitioned by the designed mechanism, thus supporting partitioning decisions. Last, its expected number of search stages can be used as an upper bound for the time it will take a sub-coalition to terminate its search (from the time of its creation), thus supporting interaction initiation decisions of other sub-coalitions.

Finally, it is notable that the model is based on several assumptions originated from classical search theory that should be carefully considered. For example, the agents are not constrained in their ability to partition their cooperation, the distribution of opportunities in the environment is known or can be learned and opportunities are explored sequentially, in a random order. While these assumptions are often used in search theory, the applicability of the proposed search model can be significantly improved if they can be overcome.

References

1. Y. Bakos. Reducing buyer search costs: Implications for electronic marketplaces. *Mgmt. Science*, 43(12) (1997) 1676–92.
2. M. Baye, J. Morgan, and P. Scholten. Temporal price dispersion: Evidence from an online consumer electronics market. *Journal of Interactive Marketing*, 18(4):101–115, 2004.
3. M. Baye, J. Morgan, and P. Scholten. Persistent price dispersion in online markets. In D. Jansen, editor, *The New Economy and Beyond*. Edward Elgar Press, Northampton, MA, 2006.
4. J. Benhabib and C. Bull. Job search: The choice of intensity. *J. of Political Economy*, 91(5) (1983) 747–764.
5. S. Breban and J. Vassileva. Long-term coalitions for the electronic marketplace. In *B. Spencer, ed., Proceedings of E-Commerce Applications Workshop, Canadian AI Conference, Ottawa*, 2001.
6. E. Brynjolfsson, E. Hu, and M. Smith. Consumer surplus in the digital economy: Estimating the value of increased product variety at online bookseller. *Management Science*, 49(11):1580–1596, 2003.
7. K. Burdett and D. A. Malueg. The theory of search for several goods. *Journal of Economic Theory*, 24(3) (1981) 362–376.
8. J. A. Carlson and R. P. McAfee. Joint search for several goods. *Journal of Economic Theory*, 32(2) (1984) 337–345.
9. S. Choi and J. Liu. Optimal time-constrained trading strategies for autonomous agents. In *Proc. of MAMA'2000, Wollongong, Australia*, 2000.
10. K. Clay, R. Krishnan, E. Wolff, and D. Fernandes. Retail strategies on the web: Price and non-price competition in the online book industry. *Journal of Industrial Economics*, 50:351–367, 2002.
11. E. David, R. Azoulay-Schwartz, and S. Kraus. Protocols and strategies for automated multi-attribute auctions. In *AAMAS '02: Proceedings of the first international joint conference on Autonomous agents and multiagent systems, Bologna, Italy*, 2002, pp. 77–85.
12. M. Dias. *TraderBots: A New Paradigm for Robust and Efficient Multirobot Coordination in Dynamic Environments*. Ph.D. thesis, Robotics Institute, CMU, 2004.
13. S. Gal, M. Landsberger, and B. Levykson. A compound strategy for search in the labor market. *International Economic Review*, 22(3) (1981) 597–608.
14. J.R.J. Gatti. Multi-commodity consumer search. *Journal of Economic Theory*, 86(2) (1999) 219–244.
15. C. V. Goldman and S. Zilberstein. Optimizing information exchange in cooperative multi-agent systems. In *AAMAS '03: Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, pages 137–144, New York, NY, USA, 2003. ACM Press.
16. P. E. Hart, N. J. Nilsson, and B. Raphael. "A formal basis for the heuristic determination of minimum cost paths". *IEEE Transactions on Systems, Science and Cybernetics*, 4(2) (1968) 100–107.
17. 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) (2002) 11–20.
18. J.P. Kahan and A. Rapoport. *Theories of coalition formation*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1984.
19. R. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. John Wiley & Sons, New York, 1976.
20. J. Kephart and A. Greenwald. Shopbot economics. *Journal of Autonomous Agents and Multi-Agent Systems*, 5(3) (2002) 255–287.
21. D. Leech. Ownership concentration and the theory of the firm: a simple-game-theoretic approach. *The Journal of Industrial Economics*, 35(3) (1987) 225–240.
22. K. Lermann and O. Shehory. Coalition formation for large scale electronic markets. In *ICMAS'2000 Boston MA*, 2000, pp. 216–222.
23. C. Li, S. Chawla, U. Rajan, and K. Sycara. Mechanisms for coalition formation and cost sharing in an electronic marketplace. In *EC03, Pittsburgh, Pennsylvania, USA*, 2003, pp. 68 – 77.
24. S. Lippman and J. McCall. The economics of job search: A survey. *Economic Inquiry*, 14(2-3) (1976) 155–189.
25. P. Maes, R. H. Guttman, and A. Moukas. Agents that buy and sell. *Commun. ACM*, 42(3)(1999) 81–91.
26. E. Manisterski. *Protocols and Strategies for Agents Teamwork*. PhD thesis, Bar Ilan University, Ramat-Gan, Israel, 2007.
27. E. Manisterski, D. Sarne, and S. Kraus. Enhancing mas cooperative search through coalition partitioning. In *Manuela M. Veloso, editor, IJCAI*, pages 1415–1421, 2007.
28. E. Manisterski, D. Sarne, and S. Kraus. Cooperative Search with Concurrent Interactions. In *Journal of Artificial Intelligence Research (JAIR)*, 32 (2008) 1–36.
29. J. McMillan and M. Rothschild. Search. In *Robert J. Aumann and Sergiu Hart, editors, Handbook of Game Theory with Economic Applications*, 1994, pp. 905–927.
30. P. Michiardi and R. Molva. Analysis of coalition formation and cooperation strategies in mobile ad hoc networks. *Ad Hoc Networks*, 3:2005, 2005.

31. P. Morgan. Search and optimal sample size. *Review of Economic Studies*, 50(4) (1983) 659–675.
32. P. Morgan and R. Manning. Optimal search. *Econometrica*, 53(4) (1985) 923–944.
33. E. Ogston, B. Overeinder, M. Van Steen, and F. Brazier. A method for decentralized clustering in large multi-agent systems. In *AAMAS03, Melbourne, Australia*, 2003, pp. 789–796.
34. T. Rahwan, S. D. Ramchurn, V. D. Dang, A. Giovannucci, and N. R. Jennings. Anytime optimal coalition structure generation. In *22nd Conference on Artificial Intelligence (AAAI)*, pages 1184–1190, February 2007.
35. M. Rothschild. Searching for the lowest price when the distribution of prices is unknown. *Journal of Political Economy*, 82(4) (1974) 689–711.
36. N. Ohta, A. Iwasaki, M. Yokoo, K. Maruono, V. Conitzer, and T. Sandholm. A compact representation scheme for coalitional games in open anonymous environments. In *AAAI06 Boston, MA, USA*, 2006.
37. D. Sarne and S. Kraus. The search for coalition formation in costly environments. In *In Proc. of CIA*, pages 117–136, 2003.
38. D. Sarne and S. Kraus. Cooperative exploration in the electronic marketplace. In *Proc. of AAAI 2005, Pittsburgh, Pennsylvania, USA*, 2005, pp. 158–163.
39. O. Shehory and S. Kraus. Methods for task allocation via agent coalition formation. *Artificial Intelligence*, 101(1-2) (1998) 165–200.
40. G. Stigler. The economics of information. *Journal of Political Economy*, 69(3) (1961) 213–225.
41. T. Sandholm, K. Larson, M. R. Andersson, O. Shehory, and F. Tohme. Coalition structure generation with worst case guarantees. *Artificial Intelligence*, 111(1-2) (1999) 209–238.
42. S. Talukdar, L. Baerentzen, A. Gove, and P. S. de Souza. Asynchronous teams: Cooperation schemes for autonomous agents. *Heuristics*, 4(4) (1998) 295–321.
43. N. Tsvetovat, K. Sycara, Y. Chen, and J. Ying. Customer coalitions in electronic markets. In *Conference on Electronic Commerce, Minnesota, MN*, 2000, pp. 121–138.
44. 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, Montreal, CA*, 2001, pp. 576–583.
45. D. Yun Yeh. A dynamic programming approach to the complete set partitioning problem. *BIT*, 26(4):467–474, 1986.

A Summary of Notations

Notation	Meaning
$B = (B_1, B_2, \dots, B_{ B })$	The set of the attributes defining any of the potentially available opportunities in the market, where each attribute B_i can be assigned a value from the finite set $(b_{min}^i, \dots, b_{max}^i)$.
O	The space of potential opportunity types the coalition may encounter.
A_g	The coalition of agents.
$U_j(\mathbf{o})$	Agent A_j 's utility from type \mathbf{o} opportunity.
$c(n)$	The search cost associated with having a coalition of size n .
θ_{known}	The set of opportunities known to the coalition at a given stage of its search.
S_{A_g}	The set of all possible states of coalition A_g .
$state(A_g, \theta)$	The state of coalition A_g acquainted with a set θ of known opportunities.
$V^*(A_g, s)$	The expected utility of a coalition A_g at state s from this point onwards when using its optimal strategy.
$V_l(A_g, s)$	The immediate utility of coalition A_g if it terminates the search at state s .
$V_r(A_g, s)$	The expected utility that coalition A_g can obtain if it resumes the search at state s .
$O_{improve}^s$	The set of opportunities that changes the coalition's current state.
O_{stay}^s	The set of opportunities that does not change the coalition's current state.
M_{A_g}	The set of all possible partitions of coalition A_g .
$V_p(A_g, s)$	The expected utility of the partitioned coalition A_g , where the selected partition will be the one yielding the maximum expected utility.
$V_{one}(s, a_j)$	The additional expected utility (without incorporating the search cost) that agent $a_j \in A_g$ obtains from terminating its search after conducting one additional search stage rather than terminating the search in the current state s it is in.
$A_{order} = (a'_1, \dots, a'_{ A_g })$	The list of agents in A_g sorted in a descending order according to their V_{one} values.
$resume(A_g, s)$	The agents that resume the search according to algorithm 1, where coalition A_g reaches state s .
$terminte(A_g, s)$	The agents that terminate the search according to algorithm 1, where coalition A_g reaches state s .