# Constraining Information Sharing to Improve Cooperative Information Gathering

Igor Rochlin
Department of Computer Science
Bar-Ilan University
Ramat-Gan, Israel
igor.rochlin@gmail.co.il

David Sarne
Department of Computer Science
Bar-Ilan University
Ramat-Gan, Israel
sarned@cs.biu.ac.il

## ABSTRACT

This paper considers the problem of cooperation between self-interested agents in acquiring better information regarding the nature of the different options and opportunities available to them. By sharing individual findings with others, the agents can potentially achieve a substantial improvement in overall and individual expected benefit. Alas, when it comes to self-interested agents, it is well known that equilibrium considerations often dictate solutions that are far from the fully cooperative ones, hence the agents do not manage to fully exploit the potential benefits encapsulated in such cooperation. In this paper we introduce, analyze and demonstrate the benefit of two methods aiming to improve cooperative information gathering. Common to all two that they constrain and limit the information sharing process. Nevertheless, the decrease in benefit due to the limited sharing is outweighed by the resulting substantial improvement in the equilibrium individual information gathering strategies. The equilibrium analysis that is given in the paper, which, in itself, is an important contribution to the study of cooperation between self-interested agents, enables demonstrating that for a wide range of settings with the use of the two methods all agents end up with an improved individual expected benefit.

## Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent Systems*

## General Terms

Economics, Algorithms

## Keywords

Multi-Agent Exploration, Self-Interested Agents, Cooperation, Teamwork, Economically-Motivated Agents

## 1. INTRODUCTION

In many settings agents can benefit from cooperating in information gathering [8, 17, 20, 16]. For example, consider two travel agents, from the same city, that plan to participate in an international tourism conference, taking place in a highly traveled destination. There are many airlines offering flights to nearby destinations, each setting a price according to various external factors such as seat availability and agreements it has with its airlines partners. Similarly, depending on the airport of arrival, one can get to the conference by train, bus, ferry, taxi or any combination of these for different segments of the trip. Each of these means of transportation may be characterized by a different availability and fare, depending, for example, on the time of the day when it is required. Checking the feasibility and cost of the different alternatives for traveling to the conference, thus, potentially involves several time consuming activities, such as checking locations on the map and checking the companies' web-sites for routes, timetables, fares and availability, and thus incurs some "opportunity cost". Since both agents can benefit from the information each of them gathers regarding the different options for getting to the conference, they have a strong incentive to share their findings, i.e., execute the information gathering process (hereafter denoted IGP) *cooperatively*.

Cooperative information gathering is used in many real-life applications of different domains. For example, consider two friends, both interested in buying a big TV screen. The friends can visit the shopping mall, together, while each of them checks offers in different stores, and eventually they meet and share their findings. Alternatively, consider an oil drilling company sending multiple agents to explore possible drilling sites, in order to develop the best site discovered. Similarly when looking to fill-in a position, HR personnel can interview candidates in parallel and recruit the best candidate found.

The benefits of multi-agent cooperative information gathering are twofold. First, since each alternative (hereafter termed "opportunity") reviewed can benefit many agents, the relative cost of information gathering is reduced, while the overall welfare increases. Secondly, the task can potentially be divided according to the expertise of the different agents, if such expertise exists.[1]

Cooperative information gathering can be seen as a type of a public goods game, where all agents contribute by their individual IGP and the collective result influences the welfare of all of them. In public goods games, in general, inefficiencies in private giving commonly occur whenever the agents are self-interested [11, 10]. Similarly, it has been shown that cooperative information gathering, carried out by self-interested agents, does not result in the amount of cooperation as in the optimal fully cooperative case [20].

In this paper we propose two cooperative information gathering methods that can increase the individual benefit of all participating agents. The methods differ by the constraints they put on the information sharing process (henceforth denoted ISP). The first, de-

[1]For buyers' cooperation, the agents can also benefit from a volume discount through their cooperation; however this property holds only for that specific domain.

**Appears in:** *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), Ito, Jonker, Gini, and Shehory (editors), May 5–9, 2014, Paris, France.*

noted “Enforced probabilistic information sharing” prevents individual agents from taking part in the ISP according to some probabilistic function. The second, denoted “Cost filtered information sharing”, introduces some cost for taking part in the ISP (where the proceeds are wasted and are not returned to the agents) and allows agents to choose whether to take part in the ISP or not. Each of these methods may seem counter intuitive, because the absence of some of the agents in the ISP is harmful to all agents. Yet, in many settings, the use of these methods can be highly beneficial. This is because of the paradox embedded in the ISP option - while the sharing of information benefits all agents, the very fact that all information gathered is going to be shared, discourages agents from investing much resources in their individual IGP. Therefore, with the use of these methods the individual benefit of each agent from taking part in the ISP decreases, however the IGP carried out by that agent individually becomes more efficient. Therefore, by intelligently managing the tradeoff between the two, a more beneficial equilibrium can be achieved, which improves both the overall and individual benefits.

The paper provides a comprehensive analysis of the individual information gathering strategies used by the agents, given the strategy of others, under the different methods. For the Enforced probabilistic method the agents’ individual strategy is proven to be similar in structure to the one used with the standard cooperative information gathering method — the agent will resume information gathering as long as the best value obtained so far is lower than some reservation value (a threshold), regardless of how much more information can potentially be gathered. For the Cost filtered method, the individual strategies are proven to be based on a single reservation value for determining the benefit in additional information gathering and for deciding whether to take part in the ISP. This allows the characterization of the resulting Bayesian Nash equilibria. Using synthetic environments, we numerically demonstrate that all two methods result in substantial improvement to each of the agents’ individual expected benefit for a wide range of settings.

The results contribute to the advancement of theories of cooperation in MAS. As discussed later in the paper, the methods can be easily applied and their use can benefit both individuals planning to engage in cooperative information gathering and MAS designers of systems where cooperative information gathering is likely to take place.

In the following section we review related work, in particular emphasizing models of cooperative information gathering. In Section 4 we describe the two proposed mechanisms, present their equilibrium analysis and supplies numerical examples for the benefit that can be achieved when using them. Discussion, conclusions and directions for future research are given in Section 5.

## 2. RELATED WORK

The model analyzed in this paper is based on two important concepts: multi-agent cooperation and costly information gathering. Multi-agent cooperation has been shown to be widely effective for better achieving agents’ individual goals or improve their performance measures, especially when there are differences in the agents’ capabilities, knowledge and resources or when an agent is incapable of completing a task by itself [35, 29, 9]. It is also the main driving force behind many coalition formation models in the area of cooperative game theory and MAS [33]. Yet, the majority of cooperation and coalition formation MAS-related research tends to focus on the way coalitions are formed and consequently concerns issues such as the optimal division of agents into disjoint exhaustive coalitions, division of coalition payoffs and enforcement methods for interaction protocols. To the best of our knowledge, no prior work on multi-agent cooperation has considered the cooperation problem of a group of self-interested agents in costly information gathering settings where findings can benefit all agents.

Group-based cooperation of self-interested agents can also be found in public goods games and allocation games in general [2, 22, 11, 10]. Common to these games is that according to their equilibrium each agent individually should opt out of the cooperation as soon as possible or invest the minimum allowed. Therefore the research of cooperation in this domain is limited to repeated games [32] or settings with bounded-rational participants (e.g., people) for which cooperation to some extent is commonly exhibited [23]. The main difference between public goods allocation games and our work is the complexity of the settings used. In our settings there is much room for individual information gathering, to some extent, even if all others are “free riders”. Moreover, with the simplistic settings used in public goods games, information sharing and the way it is handled, which is the essence of our work, is irrelevant.

The second concept upon which this paper relies, i.e. costly information gathering, is of great importance whenever there is no central source that can supply an agent full immediate reliable information on the environment and the state of the other agents. The introduction of search costs into MAS models leads to a more realistic description of these environments. This is because agents are typically required to invest/consume some of their resources in order to obtain information concerning opportunities available in their environment [30].

Optimal strategies for settings where individuals need to search for an applicable opportunity when information gathering is costly have been widely studied, prompting several literature reviews [34, 21, 15]. These models, which are often termed “costly search” models or “economic search” models have been developed to the point where their total contribution is referred to as “search theory”. Within this line of work, many cooperative information gathering models have been studied, extending the theories to multi-agent (or multi-goal) environments. Examples include, among others, an attempt to purchase several commodities while facing imperfect information concerning prices or operating several robots in order to evaluate opportunities in different locations. These works differ from ours either in that they consider fully cooperative agents that attempt to maximize the overall utility [31, 14, 5, 6, 19], and thus lack any equilibrium considerations, or they assume that any agent’s IGP is constrained by the findings of the other agents, rather than augmented/improved by such findings as in our case [27]. Consequently they constitute substantially different equilibrium strategies. Models that do consider cooperative information gathering, which rely on assumptions similar to ours (e.g., [20, 16]), focus primarily on the extraction of the equilibrium strategies and investigate the influence of the different model parameters on the agents’ performance in equilibrium. None of these works, however, suggested methods for improving the cooperative information gathering in such settings, of the kind that we suggest and analyze in this paper.

More broadly, our problem can be seen as part of the field of planning under uncertainty, hence it is related to Markov decision processes (MDP) [3, 26] and decentralized Markov decision processes [4]. In these models the goal is to maximize the expected cumulative reward, which is also the objective in our case. Alas, the use of MDPs in our case is complicated by the continuous nature of the value probability distribution functions. More importantly, our analysis and proofs result in threshold-based (or interval-based) solutions which are both simpler in terms of strategy and state representation and can be derived with a substantially lesser complexity compared to solving as MDPs.

Finally, we note that the non-intuitive findings whereby methods that essentially limit information sharing and cooperation actually have a positive impact in the self-interested case follows, in spirit, earlier results in other settings. In particular, ones in which it has been shown that so-called "inefficiencies" can increase market performance, under certain circumstances. For example, in transportation economics, e.g., congestion games, taxation can change the equilibrium to a more desirable one [25, 24, 13]. Similarly, taxes can facilitate more desirable equilibria in Boolean games [12] and in centralized matching schemes [1]. In this work we show that a somewhat similar phenomenon also occurs in the context of cooperative information gathering, though the model and analysis are, of course, totally different from the above mentioned.

# 3. THE MODEL

The model considers a set $K = \{A_1, ..., A_k\}$ of fully-rational self-interested agents. Each of the agents needs to gather information pertaining to the value (e.g., benefit) of different opportunities to which it has access and eventually choose one. While the values of the different opportunities are a priori unknown, the agent is acquainted with the distribution of opportunity values. Information is gathered for one opportunity at a time and due to the resource consuming nature of the process it is considered costly in the sense that revealing the value of an opportunity incurs a fixed cost. The individual information gathering problem, as just defined, is standard and follows the assumptions commonly used in literature [7, 17, 16, 28]. The model allows agents to differ in their information gathering competence and in the set of opportunities they can potentially access, hence the cost of revealing the value of an opportunity, the distribution of opportunity values and the number of opportunities accessible to an agent, are modeled in the agent level, using $c_i$, $f_i(x)$ and $n_i$ to denote the three, respectively.

In settings where all opportunities are applicable to all agents the agents have an incentive to cooperate in information gathering in the sense that all individual findings are eventually shared with all others. While there are many ways to share the information, the focus of this paper, as with prior models of cooperative information gathering (e.g., [20]), is on setups where the ISP takes place at some pre-specified time, after all agents have completed their individual IGPs and each needs to decide on the opportunity it chooses. As in prior models of cooperative information gathering, we also assume that: (a) the agents are truthful in the sense that they always report the true values they obtain; (b) in the absence of any finding agent $A_i$'s utility is $v_0^i$; and (c) either the opportunities each agent can check are unique or the agents can a priori divide the opportunities among them such that each will be assigned a different set. It is assumed that information gathering costs and opportunity values are additive and each agent $A_i$ is interested in maximizing its expected benefit, denoted $EB_i$. The benefit of an agent is therefore the best value obtained by the group minus the costs accumulated individually along the agent's individual IGP.

The cooperative information gathering model as detailed above can be found in full or with some insignificant variations in prior literature [16, 20, 14, 6].[2] The truthfulness assumption is commonly justified by a substantial potential reputation loss, and is easily enforceable using fines. The choice of sharing findings at the end of the individual IGPs is mostly natural and commonly used in real life. More importantly, the alternative of sharing information throughout the process has a major setback in the sense that each individual agent finds information sharing to be beneficial only when it is on the receiving end, i.e., it is the one being informed that a "favorable" opportunity was found; when it is on the reporting end, the agent loses from such communication since the report can potentially encourage the other agents to terminate their individual information gathering. On the other hand sharing the information after concluding the individual IGPs is always beneficial for the agent as it gains more information, and at the same time the information it discloses does not affect the behavior of others thereafter since they also have already concluded their IGPs.

Taking the travel agents example, the opportunities represent different alternatives for reaching the conference location and their value is their total cost. The information gathering cost is the agents' cost of the time exploring alternatives. The goal of each agent is to minimize her expected expense, defined as the cost of the best alternative found by the two plus the cost of the time spent individually reviewing different alternatives. Similarly, the model can be mapped to all the other applications mentioned in the introduction.

# 4. ANALYSIS

This section is divided according to the two cooperative information gathering enhancing methods suggested in the paper. For each method, we first develop the individual expected-benefit-maximizing (optimal) information gathering strategy of an agent taking part in the process, as the best response to the other agents' strategies. Then, we show how the collective behavior is derived and extract the equilibrium set of strategies. Since the findings of the agents are a priori uncertain, we use the Bayesian Nash Equilibria concept. Finally, we demonstrate how the expected individual benefit of all agents improves, when the method is used, compared to the standard cooperative IGP. For this purpose we use a simplistic synthetic setting where both the agents and the opportunities available to them are homogeneous. The use of the homogeneous setting is more tractable numerically and enables demonstration of the main phenomena without substantial overhead. Specifically, we use a setting where all opportunities available to all agents share the same information gathering cost and probability distribution function (uniform, between 0 and 1), denoted $c$ and $f(y)$ respectively. We stress that even though such a setting is standard in costly information gathering literature [21, 18], its use in our case is merely for illustration purposes and all the results concerning individual strategies and equilibrium structure that are given in this paper are based on formal theoretical proofs.

## 4.1 Enforced Probabilistic Information Sharing

In this method each agent $A_i$ is a priori assigned some probability $P_i^{IS}$ which is used, after it has completed its individual IGP, to determine whether it is allowed to take part in the ISP. This method requires some enforcement since once the individual IGP is completed, agents obviously will benefit from taking part in the ISP, as it does not incur any cost and at the same time can improve their best finding. This issue and possible enforcement means are discussed in Section 5.

An agent's state throughout its individual IGP is represented by the subset of opportunities on which it has already gathered information, and their associated values, and consequently the remaining opportunities for which the values are still unknown. An agent's strategy is thus the mapping from a world state to a choice $\{resume, terminate\}$ where $resume$ suggests that the agent needs to gather information about an additional opportunity (a random one, since all opportunities available to a given agent are a priori alike) and $terminate$ means that the agent needs to proceed to the

[2] While some model variants consider the task to be executed by a representative agent, acting on behalf of the group, the essence of gathering costly information and trading-off costs and benefit is the same.

ISP. Theorem 1 proves that the state representation in this case can be compacted to the best value found so far (including the fallback $v_0^i$), $v$, and that the optimal strategy can be represented in terms of a single reservation value.

THEOREM 1. *Given the probability distribution function of the maximal value obtained by all other agents that take part in the ISP, denoted $\bar{f}_i(x)$, agent $A_i$'s optimal individual information gathering strategy is to set a reservation value $r_i$, where $r_i$ is the solution to:*

$$c_i = P_i^{IS} \cdot \int_{y=r_i}^{\infty} f_i(y) \int_{x=-\infty}^{\infty} (max(y,x) - max(r_i,x)) \bar{f}_i(x) dx dy \tag{1}$$
$$+ (1 - P_i^{IS}) \cdot \int_{y=r_i}^{\infty} (y - r_i) f_i(y) dy$$

*The agent should always choose to gather information on an additional opportunity (if one is available) if the best value obtained so far is below $r_i$ and otherwise it should proceed to ISP.*

PROOF. We show that each agent's strategy can be reduced to the mapping $S(v,j) \rightarrow \{resume, terminate\}$, where $v$ is the best value obtained so far and $j$ is the number of opportunities for which the values have already been obtained. Since the benefit in further information gathering when in state $(v,j)$ increases as $v$ decreases, and the cost of doing so is not affected by $v$, the optimal strategy is necessarily reservation value-based. Therefore, for each given number of opportunities $j$, the optimal individual information gathering strategy of agent $A_i$ can be characterized by the reservation value $r_i^j$ such that the agent should resume information gathering if the best value obtained so far is below $r_i^j$ and otherwise terminate its IGP.

The remainder of the proof is inductive, showing that the reservation value used by each agent remains stationary along its IGP and is calculated according to Equation 1. For the case of $j = n_i - 1$ the agent should gather information regarding the last opportunity if and only if:

$$P_i^{IS} \cdot \int_{z=-\infty}^{\infty} \max(v,z) \bar{f}_i(z) dz + (1 - P_i^{IS}) \cdot v < \tag{2}$$
$$P_i^{IS} \cdot \int_{y=-\infty}^{\infty} f_i(y) \int_{z=-\infty}^{\infty} (\max(y,v,z) \bar{f}_i(z)) dz dy$$
$$+ (1 - P_i^{IS}) \cdot \int_{y=-\infty}^{\infty} max(y,v) f_i(y) dy - c_i$$

where $v$ is the best value obtained so far. The left hand side of the equation captures the expected benefit if the individual IGP is terminated and the right hand side captures the expected benefit if information is gathered for the last opportunity. Both terms distinguish between the case where agent $A_i$ participates in the ISP i.e., with a probability $P_i^{IS}$, and when it is not allowed to. Using several mathematical manipulations we obtain that the $v$ values for which (2) holds are those lesser or equal to the $r_i$ value that satisfies (1). Therefore the theorem holds for $j = n_i - 1$.

Now assume the same $r_i$ (according to (1)) holds for any $j' > j$, for some $j$, and consider the agents' decision regarding gathering information on one more opportunity, if the best value obtained so far is $v$ and the number of opportunities for which the values were already obtained is $j$. If $v > r_i$ and the agent gather information on one additional opportunity, then regardless of the value obtained next it will definitely terminate its individual IGP thereafter (as it already has a value greater than $r_i$ and according to the induction assumption the optimal strategy thereafter is the reservation value $r_i$). Therefore the benefit obtained from further information gathering is given by:

$$P_i^{IS} \cdot \int_{y=-\infty}^{\infty} f_i(y) \int_{z=-\infty}^{\infty} (\max(y,z) - \max(v,z) \bar{f}_i(z)) dz dy \tag{3}$$
$$+ (1 - P_i^{IS}) \cdot \int_{y=-\infty}^{\infty} (y - v) f_i(y) dy - c_i$$

Alas, since the latter term decreases as $v$ increases, and obtains zero for $v = r_i$ (according to (1)), then since $v > r_i$ the term obtains a negative value, hence additional information gathering cannot be the preferred choice.

Similarly, consider the case where $v < r_i$ for $j$ and the agent chooses not to gather additional information. Here the expected benefit from resuming information gathering is necessarily greater than if resuming in state $(v, j' > j)$. However, according to the induction assumption the agent should resume information gathering in state $(v, j' > j)$, leading to a contradiction. Therefore, the reservation value for $j$ is calculated, once again, according to (1). □

Theorem 1 specifies the optimal strategy of an agent given the strategy of others. The solution of a set of $k$ equations similar to (1), one for each agent $A_i$, will supply us with a set of pure equilibria of the form $\{r_i | 1 \leq i \leq k\}$ if any exist. A mixed equilibrium in our case is defined by a probability $p_i(v,j)$ assigned to each state $(v,j)$, defining whether the agent will resume or terminate information gathering in that state. This may seem infeasible to extract, based on the infinite number of states (as value distributions are continuous). Nevertheless, in order for such a solution to hold, the agent's expected benefit from both actions (resume and terminate information gathering when in that state) must be equal. Based on the optimality of the reservation-value rule, this can hold only for states where the value $v$ equals $r_i$ as calculated according to (1). However, due to the continuous nature of $v$, the probability of actually reaching states that satisfy the above condition is zero, thus assigning such probabilities will have no effect on the other agents. The only exception for the above is the agent's strategy at the beginning of its individual IGP. Here, the state is a priori known to be $(v_0^i, 0)$ hence adding some probability for actually gathering information on one opportunity and then continuing according to $r_i$ (or otherwise going straight to ISP due to the indifference to resuming or terminating information gathering) will have an actual effect on the others. Consequently, a mixed Bayesian Nash equilibrium for our problem is of the form:

$$\{(p_i, r_i) | 1 \leq i \leq k\}$$

where $p_i$ is the probability that agent $A_i$ will initiate its individual IGP ($0 \leq p_i \leq 1$) and $r_i$ is the reservation value to be used by the agent.

Now that the individual strategy in equilibrium has been defined in its complete form (i.e., including the probabilistic aspect), we can formulate $\bar{f}_i(x)$ (the probability distribution function of the maximal value obtained along the IGP of all other agents that will take part in the ISP). For this purpose we make use of the probability that the maximum value that will be found by all the agents that take part in the ISP, except $A_i$, will be smaller than or equal to $x$, denoted $\bar{F}_i(x)$. The calculation of $\bar{F}_i(x)$ makes use of the probability that the maximum value obtained along the IGP of an agent $A_j$ (that chooses to engage in IGP and uses $r_j$), is less than

$x$, denoted $F_j^{return}(x)$, calculated according to:

$$F_i^{return}(x) = \begin{cases} F_i(x)^{n_i} & x < r_i \\ F_i(r_i)^{n_i} + \frac{1-F_i(r_i)^{n_i}}{1-F_i(r_i)}(F_i(x)-F_i(r_i)) & x \geq r_i \end{cases} \quad (4)$$

For the case where $x < r_i$ the value of all $n_i$ opportunities must result in a value below $x$. When $x \geq r_i$ there are two possible scenarios. The first is where all $n_i$ opportunities result in a value below the reservation value $r_i$, i.e., with a $F_i(r_i)^{n_i}$ probability. The second, is where the information gathering terminates right after revealing value $y$ at the $j$th opportunity such that $r_i < y < x$ (otherwise, if $y < r_i$ the information gathering should resume) and all the former $j-1$ values obtained are smaller than $r_i$ (otherwise the $j$th opportunity is not reached). The probability of the latter case occurring (summing over all values of $j \leq n_i$) can be calculated using the geometric series:

$$\sum_{j=1}^{n_i}(F_i(x) - F_i(r_i))F_i(r_i)^{j-1} = \frac{1-F_i(r_i)^{n_i}}{1-F_i(r_i)}(F_i(x) - F_i(r_i))$$

The probability distribution function of the maximum obtained throughout agent $A_i$'s IGP, denoted $f_i^{return}(x)$, is by definition, the first derivative of $F_i^{return}(x)$:

$$f_i^{return}(x) = \frac{d(F_i^{return}(x))}{dx}$$

Thus, we can now formulate the probability that the maximum value that will be found by all the agents taking part in ISP, except $A_i$, will be smaller than or equal to $x$, $\bar{F}_i(x)$:

$$\bar{F}_i(x) = \prod_{A_j \in K \wedge j \neq i}(P_j^{IS}(p_j F_j^{return}(x) + (1-p_j)) + (1-P_j^{IS}))$$

The probability distribution function $\bar{f}_i(x)$ is the derivative of $\bar{F}_i(x)$:

$$\bar{f}_i(x) = \frac{d\bar{F}_i(x)}{dx}$$

These enable us to calculate the expected benefit of agent $A_i$ when the other agents use the set of strategies $\{(p_i, r_i) \,|1 \leq i \leq k\}$. If agent $A_i$ chooses to engage in IGP then its expected benefit, denoted $EB_i(IGP)$, is given by:

$$\begin{aligned} EB_i(IGP) &= -c_i\frac{1 - F_i(r_i)^{n_i}}{1 - F_i(r_i)} \quad (5) \\ &+ P_i^{IS} \cdot \int_{y=-\infty}^{\infty} f_i^{return}(y) \int_{x=-\infty}^{\infty} max(v_0^i, y, x)\bar{f}_i(x)dxdy \\ &+ (1 - P_i^{IS}) \cdot \int_{y=-\infty}^{\infty} max(v_0^i, y) f_i^{return}(y)dy \end{aligned}$$

where the first term on the right hand side is the expected cost incurred throughout the IGP carried out by $A_i$, calculated as:

$$c_i\sum_{j=1}^{n_i}(F_i(r_i))^{j-1} = c_i\frac{1 - F_i(r_i)^{n_i}}{1 - F_i(r_i)},$$

as the number of opportunities on which information is gathered is a geometric random variable bounded by $n_i$, with a $1 - F_i(r_i)$ success probability. The second term is the expected maximum between the best value found by the agent itself (i.e., associated with a distribution $f_i^{return}(y)$) and the best value returned by the other agents (associated with a distribution $\bar{f}_i(x)$) if agent $A_i$ participate in the ISP (i.e., with a $P_i^{IS}$ probability). The last term is the expected "best" (i.e., maximum) opportunity-value found by the agent along its information gathering if the agent is not allowed to take part in the ISP (i.e., with a $1 - P_i^{IS}$ probability).

When the agent opts not to execute individual IGP at all, its expected benefit, denoted $EB_i(\neg IGP)$, is simply the expected value of the maximum value returned by the other agents, taking part in the ISP, if taking part by itself in the process, or otherwise $v_0^i$, i.e.:

$$\begin{aligned} EB_i(\neg IGP) = P_i^{IS} \cdot \int_{x=-\infty}^{\infty} max(v_0^i, x)\bar{f}_i(x)dx \quad (6) \\ + (1 - P_i^{IS}) \cdot v_0^i \end{aligned}$$

At this point, we have everything that is needed to formulate the equilibrium stability conditions. A set of strategies $\{(p_i, r_i)|1 \leq i \leq k\}$ will be in equilibrium only if the following conditions hold: (a) for every agent $A_i$ for which $p_i = 0$, $EB_i(IGP) \leq EB_i(\neg IGP)$; (b) for every agent $A_i$ for which $p_i = 1$, $EB_i(IGP) \geq EB_i(\neg IGP)$; and (c) for every agent $A_i$ for which $0 < p_i < 1$, $EB_i(IGP) = EB_i(\neg IGP)$. Therefore, in order to find the equilibrium, the stability of $3^k$ possible solutions of type $\{(p_i, r_i) \,|1 \leq i \leq k\}$ differing in the value each $p_i$ obtains ($p_i = 0$, $p_i = 1$ and $0 < p_i < 1$) needs to be checked. For every combination, the reservation values of the different agents and the probability $p_i$ of each agent that uses a non-pure mixed strategy (i.e., with $0 < p_i < 1$) should be calculated by solving a set of equations of type (1) (one for each agent characterized by $p_i = 0$) and $EB_i(IGP) = EB_i(\neg IGP)$ (one for every agent $A_i$ for which $0 < p_i < 1$). Once the appropriate reservation values and probabilities are obtained for a given set, the stability conditions need to be validated.

We note that there is no guarantee that an equilibrium will actually exist (either pure or mixed, since there are an infinite number of strategies). Also, there is no guarantee that if one exists there will be no other equilibria (i.e., multi-equilibria is possible). In the latter case, if there is one equilibrium that dominates the others in terms of the individual expected benefit each and every agent obtains then it will likely be the one used. Otherwise, there is no way of deciding which of the equilibria is the one to be used, and we do not include this question in the scope of the current paper.

We emphasize that the above analysis generalizes the analysis of the standard cooperative information gathering model [16, 20] in the sense that the latter is a specific case where the probability each agent will be allowed to take part in the ISP is one (i.e., $P_i^{IS} = 1$ $\forall 1 \leq i \leq k$). Furthermore, when the probability each agent will be allowed to take part in the ISP is zero, the solution obtained is the same as the one known for the single-agent information gathering problem [21] (since each agent relies solely on the values it obtains throughout its individual IGP).

Figure 1 depicts the agents' individual expected benefit as a function of the probability $P_i^{IS}$ used, for different information gathering costs ($c$). The setting used is the homogeneous setting described at the beginning of the section and the value of $P_i^{IS}$ is the same for all agents (i.e., $P_i^{IS} = P^{IS}$ $\forall i$). The other model parameters were set to: $k = 15$ and $n = 4$. As depicted in the figure, the maximum expected benefit (agent-wise, as all agents are alike in this case) is obtained when the participation of the agents in the ISP is not certain but rather determined probabilistically (i.e., $0 < P^{IS} < 1$). The typical pattern exhibited in the figure is an increase and then a decrease in the expected individual benefit as the probability $P^{IS}$ increases. This is explained by the fact that when $P^{IS} = 0$ each agent actually executes an individual IGP without any information sharing with others. As $P^{IS}$ increases, the agent relies more on other agents' findings. Thus, $p_i$ and $r_i$ become lower, which is bad for the group since everybody gains less from the participation of

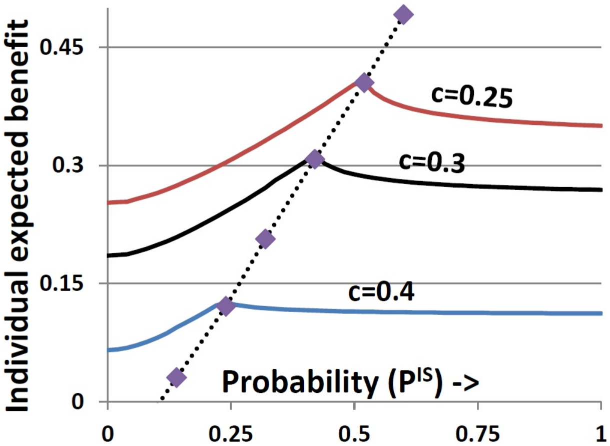


**Figure 1: Enforced Probabilistic Information Sharing - effect of $P^{IS}$ on the individual expected benefit, for different information gathering costs, $c$ in a setting: $k = 15$ and $n = 4$. The dotted curve connects the individual-expected-benefit-maximizing $P^{IS}$ values for the different $k$ and $c$ values.**

the agent in the ISP. However, at the same time the probability the agent will actually take part in the ISP increases, thus, overall, the individual expected benefit increases. Nevertheless, at some value of $P^{IS}$, the loss due to the resulting decrease in $p_i$ and $r_i$ becomes more dominant than the benefit due to the increase in the value of $P^{IS}$.

Another interesting behavior observed in Figure 1 is that the increase in information gathering costs, $c$, results in a decrease in the value of the individual-expected-benefit-maximizing ISP-participation probability (marked by the thick markers in the graph). This may seem non-intuitive since the greater the information gathering cost the greater the potential benefit that can be achieved by information sharing. Therefore further limiting information sharing in settings with high $c$ values may seem unnatural. The phenomena is explained by the fact that the positive effect of the increase in $P^{IS}$ over the participation probabilities $p_i$ and the reservation value $r_i$ used by each agent in equilibrium in this case, which are substantially poor to begin with, is greater than the loss due to the uncertain information sharing.

## 4.2 Cost Filtered Information Sharing

This method introduces a cost $c^{IS}$ an agent incurs if it chooses to take part in the ISP. The introduction of such a cost (that is not returned eventually to the agents) requires an appropriate balance in the form of some compensation to the agents for taking part in the ISP. Without such a compensation, no agent will be willing to take part in the ISP, as proven in the following theorem.

THEOREM 2. *If the agents incur some cost when taking part in the ISP, then in the absence of some compensation for taking part in the ISP, none of the agents will take part in the ISP.*

PROOF. Consider the highest value $v$ that if it turns out to be the highest value known to any of the agents, after completing its individual IGP, warrants the participation of this agent in the ISP. We show that in the absence of an appropriate compensation for the agent, such a value, $v$, cannot hold — since none of the other agents will bring a value greater than $v$ to the ISP, the agent will not gain anything from the ISP however it will incur a cost; consequently it will choose not to take part in the ISP. □

One option to compensate the agents for taking part in the ISP is to offer the agent with the best value that takes part in the ISP a compensation $B$. The amount $B$ can be collected from the agents (e.g., in equal shares) prior to the IGP, such that once collected it is considered a "sunk cost" and the agents' strategies become affected only by the chance of receiving the bonus $B$. The structure of the best response strategy of any individual agent in this case, given the strategy of others, is given in Theorem 3.

THEOREM 3. *Given the probability distribution function of the maximal value obtained by all other agents that take part in the ISP, $\bar{f}_i(x)$, agent $A_i$'s optimal individual information gathering strategy can be described by the pair $(r_i, R_i^{IS})$, where the value of $r_i$ is the solution to:*

$$c_i = \int_{y=r_i}^{\infty} (EB_i(y) - EB_i(r_i)) f_i(y) dy \qquad (7)$$

*and $R_i^{IS}$ is the set of intervals such that for any $x \in R_i^{IS}$:*

$$c^{IS} \leq \int_{y=x}^{\infty} (y - x)\bar{f}_i(y)dy + B \cdot \int_{y=-\infty}^{x} \bar{f}_i(y)dy \qquad (8)$$

*where $EB_i(v)$ is given by:*

$$EB_i(v) = \begin{cases} v & v \notin R_i^{IS} \\ -c^{IS} + B \cdot \int_{y=-\infty}^{v} \bar{f}_i(y)dy & \text{otherwise} \\ + \int_{y=-\infty}^{\infty} max(y, v)\bar{f}_i(y)dy & \end{cases} \qquad (9)$$

*The agent should resume IGP as long as the value found so far is below $r_i$, and otherwise it should terminate IGP. Upon terminating IGP (or obtaining the value of all opportunities) the agent should participate in the ISP if the best value it has found in its individual IGP is in one of the intervals of the set $R_i^{IS}$ and otherwise it should opt out from taking part in the ISP.*

PROOF. The set $R_i^{IS}$ as defined in (8) contains all the values $v$ for which the expected benefit from taking part in the ISP - calculated as the potential value improvement, $\left(\int_{y=x}^{\infty}(y - x)\bar{f}_i(y)dy\right)$ plus the expected compensation $B \cdot \int_{y=-\infty}^{x} \bar{f}_i(y)dy$, both independent of the reservation value $r_i$ used by the agent - is greater than the cost $c^{IS}$ incurred. The remainder of the proof, concerning the optimality of a reservation-value-based strategy and the correctness of (7) is the same as the one provided for Theorem 1, differing only in the way the expected benefit if resuming information gathering is calculated. □

The solution of a set of equations consisting of (7)-(9) will supply us with a set of pure equilibria of the form $\{(r_i, R_i^{IS})|1 \leq i \leq k\}$ if any exist. For the same considerations given in Section 4.1, a mixed Bayesian Nash equilibrium for this case will be of the form:

$$\{(p_i, r_i, R_i^{IS})|1 \leq i \leq k\}$$

where $p_i$ is the probability that agent $A_i$ will initiate its individual IGP ($0 \leq p_i \leq 1$), $r_i$ is the reservation value to be used by the agent and $R_i^{IS}$ is the set of intervals.

Unlike the analysis given in 4.1, here we need to distinguish between $F_i^{return}(x)$, which is the probability that the maximum value obtained by agent $A_i$'s individual IGP will be less than $x$, calculated according to (4), as before, and $F_i^{return'}(x)$, which is the probability that the maximum value provided by the agent in the ISP will be less than $x$, calculated as:

$$F_i^{return'}(x) = F_i^{return}(x) + \int_{y \geq x \wedge y \notin R_i^{IS}} f_i^{return}(y)dy$$

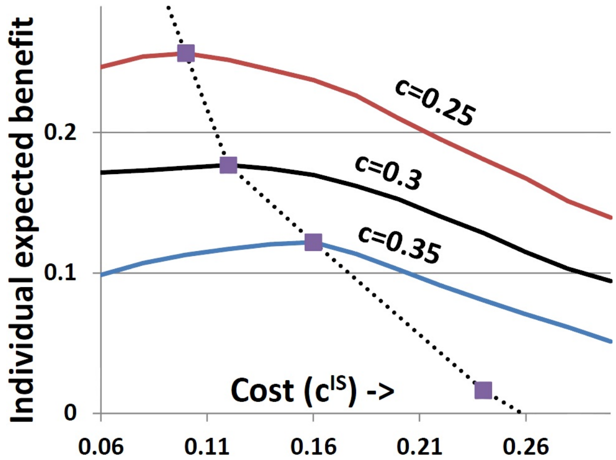


**Figure 2: Cost Filtered Information Sharing - effect of $c^{IS}$ on the individual expected benefit, for different information gathering costs, $c$ in a setting: $k = 2$, $n = 3$ and $B = 0.04$. The dotted curve connects the individual-expected-benefit-maximizing $c^{IS}$ values for the different $c$ values.**

Using $F_i^{return'}(x)$, we can now calculate the function $\bar{F}_i(x)$:

$$\bar{F}_i(x) = \prod_{A_j \in K \wedge j \neq i} (p_j F_j^{return'}(x) + (1 - p_j))$$

The probability distribution function $\bar{f}_i(x)$ is the first order derivative of $\bar{F}_i(x)$ as before. Similarly, the probability distribution function $f_i^{return}(x)$ is the first derivative of $F_i^{return}(x)$.

These enable us to calculate the expected benefit of agent $A_i$ when the other agents use the set of strategies $\{(p_j, r_j, R_j^{IS}) | 1 \leq j \leq k \wedge i \neq j\}$. If agent $A_i$ chooses to engage in IGP then its expected benefit, $EB_i(IGP)$, is given by:

$$EB_i(IGP) = -c_i \frac{1 - F_i(r_i)^{n_i}}{1 - F_i(r_i)} + \int_{y=-\infty}^{\infty} EB_i(\max(y, v_0^i)) f_i^{return}(y) dy \quad (10)$$

When the agent opts not to gather information at all, its expected benefit, $EB_i(\neg IGP)$, is simply the expected value of the maximum value returned by the other agents, if the agent chooses to take part in the ISP:

$$EB_i(\neg IGP) = EB_i(v_0^i) \quad (11)$$

The equilibrium stability conditions remain as in Section 4.1, replacing the calculation of $EB_i(IGP)$ and $EB_i(\neg IGP)$ with (10) and (11). As with the former method, there is no guarantee that an equilibrium will actually exist (either pure or mixed) and that if one exists there will be no other equilibria. Also, as with the methods presented above, the analysis of the cost filtered method generalizes the analysis of the standard cooperative information gathering model [16, 20] in the sense that the latter is a specific case where $c^{IS} = B = 0$. Similarly, when $B = 0$ the solution obtained is the same as the one known for the single-agent information gathering problem [21] (regardless of the value of $c^{IS}$, based on Theorem 2).

Figure 2 illustrates the agents' individual expected benefit as a function of the cost $c^{IS}$ used, for different information gathering costs ($c$). The setting used is the homogeneous setting described at the beginning of the section, using the parameters $k = 2$, $n = 3$ and $B = 0.04$. As can be observed from the figure, the maximum expected individual benefit is obtained when setting a substantial cost for taking part in the ISP. The typical pattern exhibited in the figure is similar to the one depicted in Figures 1, and explained by similar considerations.

# 5. DISCUSSION AND CONCLUSIONS

As demonstrated in the previous section, each of the two methods proposed and analyzed in this paper can substantially increase the benefit self-interested agents achieve through information sharing when gathering information cooperatively. Each of the two methods is based on a different restriction made on the agents' ability or willingness to take part in the ISP. Intuitively such restrictions may seem to have a negative effect on performance. Yet, since each agent gains less from the information sharing itself, it has a greater incentive to invest more resources in individual information gathering, and hence performance improves overall.

The first method require some enforcement, as the dominating strategy of the agents is to take part in the information sharing. Such enforcement is easy to achieve using simple means. For example, a designated server can be used for information sharing, enabling the establishment of communication with each agent according to the pre-defined probability. The second method enables the agents to decide whether to opt-out from information sharing hence it does not require any enforcement whatsoever. It does require, however, a means for introducing a cost for the ISP. For example, the agents can decide that each agent that decides to take part in the ISP will donate a fixed amount of money to her favorite charity.

The results suggest important inputs for the designers of markets and systems where cooperative information gathering is applicable, by enabling them to predict the strategies that will be used and the resulting system performance. These primarily facilitate the proper design of the system and the determination of what elements should and should not be included in such systems in order to achieve specific goals and promote certain behavior. In particular, the introduction of some seemingly non-beneficial elements may actually be productive. We note that the paper generally does not attempt to find the "optimal" parameter values for each method (e.g., probability and cost of taking part in the information sharing, or the payment received if the agent is associated with the "best" value), since the concept of optimality in this sense is not properly defined. Indeed in settings where there is an equilibrium solution that is preferred by all agents (e.g., in the examples given in the former section, where all agents are homogeneous) the choice of parameter values is clear. Nevertheless in general, it is possible that a certain value will be preferred by one of the agents whereas others will prefer another. In the latter case it is the role of the system designer to decide on these parameters based on her goals.

There are numerous extensions of the model that can be considered. Some of them are straightforward and require minor changes in the analysis. For example, if agents are buyers and each of them is interested in more than a single unit of the product they are searching for, the only change required in the individual strategy equations is multiplying the expense of purchasing the item by the number of items in which the agent is interested. Other extensions, while of much interest, are more complex to analyze. For example, consider a model where the agents can continuously share their findings along their individual IGPs. In this case, as discussed in Section 1, it is essential to first define the method that will provide an incentive for agents to share their findings despite the negative influence it will have in terms of discouraging others from further information gathering.

## Acknowledgments

The authors would like to thank Barbara Grosz for insightful comments that helped improve this paper substantially. This work was partially supported by ISF grant 1083/13, the Israeli Ministry of Industry and Trade under project RESCUE.

## 6. REFERENCES

[1] E. Anshelevich, S. Das, and Y. Naamad. Anarchy, stability, and utopia: Creating better matchings. In *Proceedings of the Second International Symposium on Algorithmic Game Theory*, pages 159–170, 2009.

[2] R. Aumann. On the centipede game. *Games and Economic Behavior*, 23(1):97–105, 1998.

[3] R. Bellman. A markovian decision process. *Indiana University Mathematics Journal*, 6(4):679–684, 1957.

[4] D. Bernstein, D. Givan, N. Immerman, and S. Zilberstein. The complexity of decentralized control of markov decision processes. *Mathematics of Operations Research*, 27(4):819–840, 2002.

[5] K. Burdett and D. A. Malueg. The theory of search for several goods. *Journal of Economic Theory*, 24(3):362–376, 1981.

[6] J. A. Carlson and R. P. McAfee. Joint search for several goods. *Journal of Economic Theory*, 32(2):337–345, 1984.

[7] M. Chhabra and S. Das. Learning the demand curve in posted-price digital goods auctions. In *Proceedings of the 10th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 63–70, 2011.

[8] M. Chhabra, S. Das, and D. Sarne. Expert-mediated search. In *Proceedings of the 10th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 415–422, 2011.

[9] V. Conitzer. Computing game-theoretic solutions and applications to security. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*, pages 2106–2112, 2012.

[10] S. de Jong and K. Tuyls. Human-inspired computational fairness. *Autonomous Agents and Multi-Agent Systems*, 22(1):103–126, 2011.

[11] S. de Jong, K. Tuyls, and K. Verbeeck. Fairness in multi-agent systems. *Knowledge Engineering Review*, 23(2):153–180, 2008.

[12] U. Endriss, S. Kraus, J. Lang, and M. Wooldridge. Designing incentives for boolean games. In *Proceedings of the 10th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 79–86, 2011.

[13] D. Fotakis, G. Karakostas, and S. G. Kolliopoulos. On the existence of optimal taxes for network congestion games with heterogeneous users. In *Proceedings of the Third International Symposium on Algorithmic Game Theory*, pages 162–173, 2010.

[14] J. Gatti. Multi-commodity consumer search. *Journal of Economic Theory*, 86(2):219–244, 1999.

[15] A. Grosfeld-Nir, D. Sarne, and I. Spiegler. Modeling the search for the least costly opportunity. *European Journal of Operational Research*, 197(2):667–674, 2009.

[16] N. Hazon, Y. Aumann, and S. Kraus. Collaborative multi agent physical search with probabilistic knowledge. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, pages 167–174, 2009.

[17] J. Kephart and A. Greenwald. Shopbot economics. *Journal of Autonomous Agents and Multi-Agent Systems*, 5(3):255–287, 2002.

[18] S. Lippman and J. McCall. The economics of job search: A survey. *Economic Inquiry*, 14(3):347–368, 1976.

[19] E. Manisterski, D. Sarne, and S. Kraus. Cooperative search with concurrent interactions. *Artificial Intelligence Research*, 32:1–36, 2008.

[20] M. Mash, I. Rochlin, and D. Sarne. Join me with the weakest partner, please. In *Proceedings of the 2012 IEEE/WIC/ACM International Conference on Intelligent Agent Technology*, pages 17–24, 2012.

[21] J. McMillan and M. Rothschild. Search. In *Handbook of Game Theory with Economic Applications*, volume 2, pages 905–927. 1 edition, 1994.

[22] R. Nagel and F. Tang. Experimental results on the centipede game in normal form: an investigation on learning. *Journal of Mathematical Psychology*, 42(2-3):356–384, 1998.

[23] N. Peled, M. Bitan, J. Keshet, and S. Kraus. Predicting human strategic decisions using facial expressions. In *Proceedings of the 23st International Joint Conference on Artificial Intelligence*, 2013.

[24] M. Penn, M. Polukarov, and M. Tennenholtz. Random order congestion games. *Mathematics of Operations Research*, 34(3):706–725, 2009.

[25] M. Penn, M. Polukarov, and M. Tennenholtz. Taxed congestion games with failures. *Annals of Mathematics and Artificial Intelligence*, 56(2):133–151, 2009.

[26] M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley-Interscience, 1994.

[27] I. Rochlin and D. Sarne. Information sharing under costly communication in joint exploration. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence*, pages 847–853, 2013.

[28] M. Rothschild. Searching for the lowest price when the distribution of prices is unknown. *Journal of Political Economy*, 82(4):689–711, 1974.

[29] W. Saad, Z. Han, M. Debbah, and A. Hjorungnes. Coalitional games for distributed collaborative spectrum sensing in cognitive radio networks. In *Proceedings of the IEEE INFOCOM*, pages 2114–2122, 2009.

[30] D. Sarne and S. Kraus. The search for coalition formation in costly environments. In *Cooperative Information Agents VII*, pages 117–136. Springer, 2003.

[31] D. Sarne, E. Manisterski, and S. Kraus. Multi-goal economic search using dynamic search structures. *Journal of Autonomous Agents and Multi-Agent Systems*, 21(1-2):204–236, 2010.

[32] R. Selten and R. Stoecker. End behavior in sequences of finite prisoner's dilemma supergames a learning theory approach. *Journal of Economic Behavior & Organization*, 7(1):47–70, 1986.

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

[34] L. Smith. Frictional matching models. *Annual Reviews in Economics*, 3(1):319–338, 2011.

[35] P. Stone and S. Kraus. To teach or not to teach? decision making under uncertainty in ad hoc teams. In *Proceedings of the Ninth International Conference on Autonomous Agents and Multiagent Systems*, pages 117–124, 2010.