Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI-17)

Nurturing Group-Beneficial Information-Gathering Behaviors Through Above-Threshold Criteria Setting

Igor Rochlin
igor.rochlin@gmail.com
School of Computer Science
College of Management
Rishon LeZion, Israel

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

Maytal Bremer
bmaytalb@gmail.com
School of Computer Science
College of Management
Rishon LeZion, Israel

Ben Grynhaus
ben.grynhaus@gmail.com
School of Computer Science
College of Management
Rishon LeZion, Israel

Abstract

This paper studies a criteria-based mechanism for nurturing and enhancing agents' group-benefiting individual efforts whenever the agents are self-interested. The idea is that only those agents that meet the criteria get to benefit from the group effort, giving an incentive to contribute even when it is otherwise individually irrational. Specifically, the paper provides a comprehensive equilibrium analysis of a threshold-based criteria mechanism for the common cooperative information gathering application, where the criteria is set such that only those whose contribution to the group is above some pre-specified threshold can benefit from the contributions of others. The analysis results in a closed form solution for the strategies to be used in equilibrium and facilitates the numerical investigation of different model properties as well as a comparison to the dual mechanism according to only an agent whose contribution is below the specified threshold gets to benefit from the contributions of others. One important contribution enabled through the analysis provided is in showing that, counter-intuitively, for some settings the use of the above-threshold criteria is outperformed by the use of the below-threshold criteria as far as collective and individual performance is concerned.

Introduction

Cooperation between agents is a highly desired goal in many multi-agent systems (MAS). By working together and helping each other, agents can better perform and meet their goals, both collectively and individually (Stone et al. 2010; Dutta and Sen 2003; Kraus, Shehory, and Taase 2003; Rochlin, Sarne, and Mash 2014). As such, group based cooperative behavior is studied in various domains and contexts such as solving complex optimization problems (Barbulescu et al. 2010), military and rescue applications (Dias and Sandholm 2004; Conitzer 2012), cognitive radio networks (Xie et al. 2010), e-business applications (Yamamoto and Sycara 2001) and many more.

Common to many of the above mentioned domains is that the agents contribute to the group through their individual efforts and the collective result influences the welfare of all of them. This situation, in which costs and resources are basically born individually although benefits are societal, poses a dilemma for the agents—when being self-interested, any agent contributes to the group effort only to the extent that it finds to be individually beneficial. Such inefficiencies in private giving preclude the group from achieving the collective (and consequently also the individual) performance encapsulated in the fully cooperative solution (de Jong and Tuyls 2011; Rochlin, Sarne, and Mash 2014; A. van Wissen et al. 2012). The phenomenon is best illustrated in social good allocation games (e.g., in the centipede game (Aumann 1998; McKelvey and Palfrey 1992; Peled et al. 2013)) and in public goods games in general. Common to these games is that, according to their equilibrium, each agent individually should opt out as soon as possible or invest the minimum allowed. For example, consider the hypothetical case where the city council decides to build a new country club and is willing to match whatever amount of money the residents contribute for this purpose. Assume there are $N > 2$ residents and the utility of each resident is the total funds raised for the country club (including the matching) divided by the number of residents using the country club, minus the individual contribution made. In this case, there is a single equilibrium according to which all residents contribute zero, which is far (both in terms of social welfare and individual utilities) from the fully cooperative solution according to which all residents equally contribute as much as possible.

One way of overcoming this problem is by setting criteria for eligibility to the fruits of the cooperative effort. The idea is to incentivize individuals to put in as much of an effort as is needed to comply with the criteria set. Taking the above country club example, assume that the city council decides that only those who contribute above some threshold V will be allowed to use the country club. With this new restriction, the equilibrium is that each resident contributes exactly V. In fact, the larger the value of V, the greater the social welfare of the residents.

Intuitively, when using a criteria-based mechanism for enhancing cooperation in MAS, the system designer should aim for a criteria that rewards those who contribute the most (or at least sufficiently) for the group effort. This way there is a greater individual benefit for the agent in making the extra effort in order to benefit also from the contributions of the others (who in turn will also try harder, due to the same considerations). Indeed, in many settings (including the country club example given above), establishing a criteria aiming to benefit those who contribute less than others (or less than some pre-specified performance level) is always detrimental. In this paper, however, we manage to show that the reverse relationship can hold as well. For this purpose we use a standard cooperative information gathering model where agents can benefit from the information gathered by others (Kephart and Greenwald 2002; Rochlin, Sarne, and Mash 2014; Hazon et al. 2013). For example, several buyers interested in the same product can check prices and terms in different online and physical stores and share their findings, HR personnel can interview candidates in parallel and recruit the best candidate found and students can jointly look for references for an assignment they receive and eventually use the best source found by any of them. The individual information gathering process is costly, in the sense that the agent needs to consume some resources in order to learn about the different opportunities, hence the underlying conflict described above.

The contributions the paper makes are twofold: First, it provides a comprehensive equilibrium analysis for the cooperative information gathering model under an "above-threshold" criteria (i.e., one that enables only agents who contribute to the group a finding better than some pre-specified threshold to become eligible for benefiting from the findings of others). As part of the analysis of this case, which is in itself a contribution to the study of cooperative information gathering, we show that the strategies used by the agents in equilibrium are different (in structure) than those used under the dual "below-threshold" criteria and provide a closed form solution for extracting the equilibrium strategies. Second, we manage to show that (at times) benefiting those who poorly contribute to the group (using the below-threshold criteria) can be the dominating mechanism. Furthermore, the paper manages to exemplify this also for settings of homogeneous agents (i.e., with equal information gathering capabilities), meaning that such dominance of a below-threshold mechanism is not always the result of differences between the capabilities of the different agents, as one might expect.

The Cooperative Information Gathering Model

We follow the cooperative information gathering model used in our prior work (Rochlin, Sarne, and Mash 2014), which can also be found in full or with some variations in other prior literature (Hazon et al. 2013; Gatti 1999; Carlson and McAfee 1984). The model considers a set $K = \{A_1, \ldots, A_k\}$ of fully-rational self-interested agents, each engaged in gathering information pertaining to the value (e.g., benefit) of different opportunities to which it has access (with no overlap between the sets of opportunities available to different agents). The uncertainty associated with the value of opportunities available to any agent A_i is modeled through a probability distribution function (p.d.f.) $f_i(x)$. The process is considered costly in the sense that revealing the value of an opportunity incurs a fixed cost, denoted c_i. The model assumes that any agent A_i is constrained by the number of opportunities it can access, denoted n_i. The agent thus needs to gather information, i.e., explore the value of some of the opportunities and eventually pick one of the values revealed (i.e., recall is permitted). This individual information gathering process (hereafter denoted IGP) is standard and widely found in literature (Kephart and Greenwald 2002; McMillan and Rothschild 1994; Tang, Smith, and Montgomery 2010; Waldeck 2008).

The model assumes that all opportunities accessible to an agent are applicable to all other agents as well, hence the agents can benefit from sharing their findings upon completing the IGP. The information sharing process (hereafter denoted ISP) is assumed to be truthful in the sense that agents always report their true findings. We also assume that each agent A_i has some fall-back value v_0^i, i.e., even if it does not become acquainted with any opportunity value (in case of not executing an individual IGP and not taking part in ISP) the agent can presumably benefit v_0^i. In order to induce effective information gathering by the individual agents, the system designer can set a criteria such that only those contributing findings that meet the criteria become eligible for taking part in the ISP. While various criteria may be applied, we deliberately limit ourselves to threshold-based ones. The threshold-based criteria can be either of type "above-threshold" (hereafter denoted AT) or "below-threshold" (denoted BT). With both types a threshold V^{IS} is set for each agent A_i and the agent is allowed to take part in the ISP only if it ends up with a value greater than the threshold (in the case of AT) or below the threshold (in the case of BT) after its individual IGP (and otherwise denied). The model assumes that all agents are symmetric in the sense that they are all a priori familiar with $f_i(x)$, c_i, v_0^i, n_i $\forall i$ and the threshold V^{IS} .

It is assumed that information gathering costs and opportunity values are additive and each agent A_i is interested in maximizing its expected profit, denoted EB_i. The profit of an agent is therefore the best value obtained by the group minus the costs accumulated individually along the agent's individual IGP.

The essence of the cooperative information gathering is thus that individuals contribute through their individual costly IGP and the collective result influences the welfare of all of them. The model applies to various real-life settings (e.g., see Table 1).

Beyond the wide applicability of the cooperative information gathering model, we find several other reasons that make it ideal for our purposes. First, in prior work we have shown that the results achieved by self-interested agents under this model are far from those that can be achieved when the agents are fully cooperative and aim to maximize the accumulated expected profit (Rochlin, Sarne, and Mash 2014). The reason is that each agent chooses to engage in IGP to the extent it finds individually beneficial, given the IGPs carried by others. Second, in other prior work we have demonstrated that the agents' individual expected profit can be improved by setting various criteria for becoming eligible to take part in the ISP rather than by enabling this to all agents as a default (Rochlin and Sarne 2015). Finally, the case of setting a BT criteria was analyzed for this model in that work, which

Application	Individual goal	Opportunity	Value	Cost	Source of uncertainty
Product acquisition	Minimize individual expense	A store selling the product	Posted price	Time spent; communication/transportation expenses	Inter-store competition and seasonal effects
Information Search (student's assignment)	Maximize individual utility (grade minus individual effort)	Information source (e.g., online, textbook, library resource)	Expected grade if this source is used	Time spent	Differences in coverage of the topic, relevance, accuracy, level of details
Choosing an oil-drilling site	Maximize oil revenues minus cost of exploratory drills	Potential drilling sites	Amount of oil found	Time and resources spent in the exploratory drills	Uncertainty regarding the amount of oil in each drilling site
R&D	Maximize the cost of production and R&D expenses	Production technology	Cost of production with a specific technology	R&D cost of specific technology	Uncertainty concerning implementation aspects of a desired technology

Table 1: The mapping of different applications to the cooperative information sharing problem.

saves us the analysis of this variant when comparing the two.

Analysis

The main advantage of the AT method analyzed in this section is that only those with a potentially substantial contribution to the group get to benefit from the fruits of the joint effort. A substantial contribution is not always the result of a substantial effort (e.g., the agent could have run into a favorable value upon gathering information at the first opportunity encountered). Still, in general, the more effort an agent makes, the greater its chance of ending up with a potentially substantial contribution to the group. In particular, an agent that opts not to engage in IGP will end up with no benefit whatsoever, as it will always be excluded from the ISP. This creates a great incentive for agents to engage in a substantial IGP and also coincides with social norms (e.g., see La Fontaine's "The Grasshopper and the Ants" fable or even in Thessalonians 3:10: "If a man will not work, he shall not eat").

Best Response Strategy

We begin with the best-response individual IGP strategy of an agent, given the IGP strategies used by the other agents in the group. This is captured by Theorem 1.

Theorem 1. *Given the information gathering strategy of the other agents, agent A_i's expected-profit-maximizing (hereafter denoted optimal) individual IGP strategy is to follow a reservation value r_i, where r_i is either the solution to:*[1]

$$c_i = \int_{y=r_i}^{\infty} f_i(y) \int_{x=-\infty}^{\infty} (\max(y,x) - \max(r_i,x))\bar{f}_i(x)dxdy \tag{1}$$

or:

$$c_i=\int_{y=r_i}^{V^{IS}}(y-r_i)f_i(y)dy+\int_{y=V^{IS}}^{\infty}f_i(y)\int_{x=-\infty}^{\infty}(\max(y,x)-r_i)\bar{f}_i(x)dxdy \tag{2}$$

where $\bar{f}_i(x)$ is the probability distribution function of the maximal value obtained by all other agents that take part in the ISP. Given the best value obtained so far, v, agent A_i should resume its IGP if $v < r_i$ or terminate the IGP otherwise. The reservation value r_i thus serves as a threshold for deciding on resuming or terminating the IGP.

[1]In case $F_i(v) = 0$ for some value $v > -\infty$, it is possible that none of the equations yields a solution, in which case the optimal reservation value to be used is $r_i = v$.

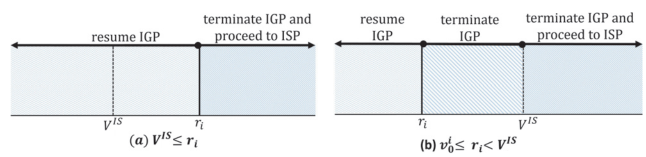

Figure 1: A schematic illustration of the optimal IGP strategy under AT.

The detailed proof is provided in the full version of the paper that is downloadable from the corresponding author's web-site. Its flow is based on first proving the reservation-value nature of the optimal strategy. Then, using an inductive proof, we show that the reservation value used by each agent remains stationary along its IGP and that it is captured by either (1) or (2).

The difference between Equations 1 and 2 is in the relationship between r_i and the threshold V^{IS} (see Figure 1). Equation 1 corresponds to the case where $r_i \geq V^{IS}$. Here, the agent will keep on gathering information on opportunities until either exhausting the set of n_i opportunities it can gather information on, or reaching a value that warrants taking part in the ISP. Equation (2) corresponds to the case where $r_i < V^{IS}$. Here, it is possible that the agent will terminate its IGP (without exhausting all available opportunities) even if the best value found so far does not enable it to take part in the ISP. Note that this single-reservation-value strategy structure is very different from the structure of the strategy known to be the best response when using a BT criteria (see the detailed analysis of this case in Rochlin and Sarne (2015)). For the latter mechanism it has been shown that the optimal IGP is based on two reservation values $r_i(j)$ and r_i^{resume} (see Figure 2). Given the best value obtained so far, v, agent A_i should resume its IGP if $v < r_i(j)$ or $V^{IS} < v \leq r_i^{resume}$ and otherwise terminate (and proceed to ISP if $v \leq V^{IS}$). Intuitively, the reservation value r_i^{resume} is used to determine if the IGP should be resumed in cases where a favorable value $v > V^{IS}$ has been found and now

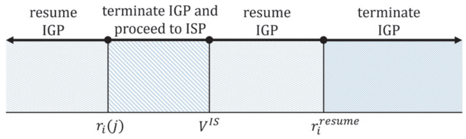

Figure 2: A schematic illustration of the optimal IGP strategy under BT criteria (taken from (Rochlin and Sarne 2015)).

the agent is "on its own". The reservation value $r_i(j)$ (which changes according to the number of opportunities remaining to the agent for gathering information) is used to determine if the IGP should be resumed when $v \leq V^{IS}$, i.e., when the agent can still potentially take part in the ISP. Resuming the IGP in the latter case can lead to better values, however at the same time can also lead to exclusion from the ISP, in which case the agent will end up on its own. The best response strategy of AT is thus different from the one used under BT both in the sense that it uses a single reservation value and that the reservation value used is stationary (i.e., does not depend on the state of the IGP).

Equilibrium

Solving the 2^k possible equation sets (where each set is of k equations, each capturing the reservation value of a different agent, either according to (1) or (2)), will provide a set of pure equilibria of the form $\{r_i|1 \leq i \leq k\}$ if any exist. A mixed equilibrium is also possible in this case, yet since the optimal IGP strategy, once engaged, was proven to be deterministic in Theorem 1, a randomization is possible only in the agent's decision whether to engage in IGP or not. Thus, a mixed 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_i, is less than (or equal to) x, denoted $F_i^{\text{return}}(x)$, calculated according to:

$$F_i^{\text{return}}(x) = \begin{cases} F_i(x)^{n_i} & x \leq r_i \\ F_i(r_i)^{n_i} + \frac{1-F_i(r_i)^{n_i}}{1-F_i(r_i)}(F_i(x) - F_i(r_i)) & x > r_i \end{cases}$$

In a case where $x \leq r_i$, the value of all n_i opportunities must result in a value below x. When $x > r_i$ 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 l^{th} opportunity such that $r_i < y < x$ (otherwise, if $y < r_i$ the information gathering should resume) and all the former $l-1$ values obtained are smaller than r_i (otherwise the l^{th} opportunity is not reached). The probability of the latter case occurring (summing over all values of $l \leq n_i$) can be calculated using the geometric series $\sum_{l=1}^{n_i}(F_i(x) - F_i(r_j))F_i(r_i)^{l-1} = \frac{1-F_i(r_i)^{n_i}}{1-F_i(r_i)}(F_i(x) - F_i(r_i))$.

Thus, we can now formulate the probability that the contribution of agent A_i to the other agents in the group that are taking part in the ISP is less than (or equal to) x, denoted $F_i^{\text{return}'}(x)$:

$$F_i^{return'}(x) = \begin{cases} F_i^{return}(V^{IS}) & x \leq V^{IS} \\ F_i^{return}(x) & V^{IS} < x \end{cases}$$

The case where $x \leq V^{IS}$ corresponds to all scenarios where agent A_i completed its IGP with a maximum value lower than the threshold V^{IS} set to it (and therefore it does not take part in the ISP and consequently its contribution to those participating in the ISP is 0, which is lower than x), i.e., with probability $F_i^{return}(V^{IS})$. The second case simply corresponds to all cases where agent A_i completed its IGP with a maximum value lower than x.

Using $F_i^{\text{return}'}(x)$, we can calculate the function $\bar{F}_i(x)$, which is the probability that the value agent A_i will gain from taking part in the ISP is less than (or equal to) x:

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

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

These enable us to calculate the expected profit of agent A_i when the other agents use the set of strategies $\{(p_i, r_i)|1 \leq i \leq k \wedge i \neq j\}$. If agent A_i chooses to engage in IGP and finds a value greater than V^{IS}, then its expected profit, denoted $EB_i(IGP)$, is given by:

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

where the first term on the right hand side corresponds to the case of failing to qualify for the ISP. The second term is the expected maximum between the best value found by the agent itself and the best value contributed by all others, corresponding to the case of qualifying for the ISP. The last term represents the expected cost incurred throughout the IGP carried out by A_i, calculated based on the expected number of values gathered (which is a geometric random variable bounded by n_i, with a $1 - F_i(r_i)$ success probability).

When the agent opts not to gather information at all, it does not participate in the ISP and therefore its expected profit, denoted $EB_i(\neg IGP)$, is simply v_0^i, i.e., $EB_i(\neg IGP) = v_0^i$.

Methods Comparison

The main difference between the AT and BT criteria is that, once an agent is engaged in IGP, the first criteria is pushing it to deepen its information gathering such that it will meet the criteria and benefit from the findings of others. With the BT criteria, agents that engage in IGP may actually lose from extending their information gathering as they may run into a value that will disqualify them from taking part in the ISP. On the other hand, with the AT criteria, the incentive to continue gathering information once meeting the criteria drops significantly, whereas with BT once an agent finds a value greater than the disqualifying threshold it has a great incentive to keep gathering information as it is "on its own".

Another important difference between the AT and BT criteria is that the first creates a very strong incentive for each agent to engage in IGP from the outset. In the BT, agents that do not engage in IGP necessarily benefit from the findings of others. The main implication is thus a plethora of mixed equilibria in which some of the agents do not engage in IGP, hoping to benefit from the findings of others. With the AT criteria, the typical equilibrium is based on pure strategies. In fact, it is enough that each agent will find the myopic profit from gathering information on a single opportunity to be positive to guarantee that all resulting equilibria will be based on pure strategies.

Proposition 1. *If* $\int_{y=v_i^0}^{\infty}(y-v_i^0)f_i(y) > c_i \ \forall i$, *the resulting equilibrium is based on pure strategies only, where all agents engage in IGP.*

Proof. Given that its profit from gathering information on one opportunity is positive, the agent will necessarily engage in IGP (which, if used according to Theorem 1, will result in at least the same profit as gathering information on only one opportunity), regardless of the strategies used by the others, as otherwise its exclusion from the ISP will prevent any profit whatsoever. The same holds true for all other agents. □

Neither of the above inherent differences between AT and BT imply any dominance relationship between the two. In fact, as we found in many of the examples we analyzed, with some settings the dominance relationship changes simply by changing the value of one of the parameters which define the specific setting. Still, the fact that with the AT criteria one is likely to expect pure equilibria suggests several advantages. For example, a pure equilibrium, is likely to result in a low variance in the expected profit of the different agents (as the number of agents engaging in information gathering is constant), compared to the case of an equilibrium based on mixed strategies.

Finally, we note that the AT criteria enables a closed form solution—the equilibrium strategies can be extracted using a set of equations, whereas with BT the solution is complicated by the need to compute the set of reservation values based on all possible states (v, j) (corresponding to the best value found and the number of remaining opportunities) which is potentially infinite, hence it requires numerical approximation.

Numerical Illustration

Figure 3 depicts the individual expected profit of the agents as a function of different parameters of the model, when using: (a) the equilibrium solution for the AT criteria; (b) the equilibrium solution for the BT criteria (according to (Rochlin and Sarne 2015)); and (c) the equilibrium solution when no criteria is applied and the agents always take part in the ISP (equivalent to setting an infinite BT criteria). The setting used considers the agents to be homogeneous in terms of their information gathering environment, i.e., sharing the same probability distribution function $f(y)$ ((i.e., $f_1(y) = .. = f_k(y) = f(y)$)), constrained to the same number of opportunities n they gather information on (i.e., $n_1 = .. = n_k = n$), sharing the same information gathering cost c (i.e., $c_1 = .. = c_k = c$) and having the same fall-back value (i.e., $v_0^1 = .. = v_0^k = v_0$). Such a setting is quite common in real-life (Rochlin, Sarne, and Mash 2014; Hajaj, Hazon, and Sarne 2015), especially when the information pertaining to different opportunities can be found online, hence the effort associated with the information gathering is quite standard (e.g., time spent navigating through a web-site). Specifically, in this example we use $f(y) = 1$ for any $0 \le y \le 1$ and $f(y) = 0$ otherwise, and a fall-back value $v_0 = 0$. The other setting parameters used are: $n = 5$ and $k = 5$ (for graph (a)); $n = 5$ and $c = 0.2$ (for graph (b)); and $k = 3$ and $c = 0.2$ (for graph (c)). The threshold used with each method is the one that maximizes the expected profit of the agents (and hence also maximizes the social welfare). We note that the patterns observed in the expected profit as a function of the different parameters (a decrease due to an increase in the information gathering cost, an increase due to an increase in the number of agents and due to the number of opportunities each agent can gather information on) are not general—we have other numerical settings that exemplify different relationships. Still, the figure reliably represents a general pattern we have observed in numerous settings that we have randomly generated: the extent of improvement achieved with the BT criteria is somehow limited and in many settings the optimal threshold is set to a relatively large value such that it has no influence on the individual IGPs compared to the case of allowing all agents take part in the ISP. The AT criteria, on the other hand, manages to affect the IGPs in all settings, and the extent of influence it achieves actually increases as the problem becomes "richer" (e.g., when the number of agents and/or number of opportunities increase). When the information gathering cost is zero, the performance with all three mechanisms is identical as the agents can only benefit from gathering more information.

While in Figure 3 the AT criteria dominate BT, this is not always the case. The following two examples illustrate that (somehow counter-intuitively) BT can, at times, dominate AT. Each example captures different types of dynamics that lead to the phenomena. The first example uses once again an homogeneous setting, this time however with two agents ($k = 2$), three available opportunities to each agent ($n = 3$), an information gathering cost of $c = 0.1$, and a uniform distribution of values (defined over the interval (0,1) as before). One important motivation for using an homogeneous set-

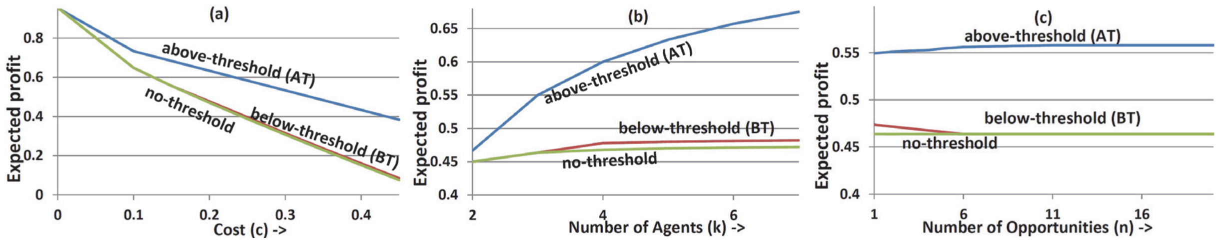

Figure 3: Expected profit of different criteria as a function of the: (a) information gathering cost c; (b) number of agents k; and (c) number of opportunities n. See the main text for the details of the settings used.

tings for this example is that the result according to which BT sometimes performs better than AT becomes even more counter-intuitive—with heterogeneous agents one may explain such phenomena by the inherent differences between the agents' individual information gathering competence. Meaning that with BT one can potentially influence some of the less competent agents to avoid the IGP and still benefit from the findings of others, leading to both better social and individual benefits. In this example there is only one equilibrium for each criteria. With BT the equilibrium is based on mixed strategies such that $p_1 = p_2 = 0.96$. The individual expected profit in this case is $EB_1 = EB_2 = 0.611$. With AT, the equilibrium is based on pure strategies, such that both agents engage in information gathering (with $r_1 = r_2 = 0.221$). The individual expected profit in this case is $EB_1 = EB_2 = 0.59$. While with the AT criteria both agents are forced to engage in IGP, in the absence of criterion each agent has an incentive to terminate its IGP after gathering information on one opportunity, due to the relatively high information gathering cost and the recognition that all agents are engaged in IGP. Therefore, in order to push the agents to extend their IGP, the optimal threshold is set to $V^{IS} = 0.22$ and consequently there is an increase in the expected IGP extent. Alas, the increase in the ISP threshold also has a downside in the sense that when the number of opportunities available is limited, the agents may frequently fail to comply with the criteria set, even if gathering information on all the opportunities available. With BT, the optimal threshold is relatively high ($V^{IS} = 0.993$) and the information sharing is taken to the greatest extent.

The second example considers an heterogeneous setting with four agents ($k = 4$). There is only one opportunity each agent can explore. The difference between the agents is in their underlying distribution function of values: while three of the agents use a uniform distribution function defined over the interval $(0, 1)$, the fourth uses a uniform distribution function defined over $(1, 2)$. The information gathering cost of all agents is similar: $c_1 = c_2 = c_3 = c_4 = 0.2$, and so is the fall-back value: $v_0^1 = v_0^2 = v_0^3 = v_0^4 = 0$. In this case both optimal criteria result in a single equilibrium based on pure strategies where with BT only the fourth agent engages in IGP and with AT all agents engage in IGP. The individual expected profit when using BT is $EB_1 = EB_2 = EB_3 = 1.5$ and $EB_4 = 1.3$, resulting in an average expected profit of 1.45. With AT the expected individual profit is $EB_1 = EB_2 = EB_3 = EB_4 = 1.3$ and so is the average expected profit. Here, the fact that with the AT criteria all agents are forced to engage in IGP actually serves them badly because the resulting cooperative IGP is highly inefficient. The fully cooperative strategy in this case is to have only one of the agents engage in IGP and have the other benefit from the finding and this is exactly the solution achieved with the BT criteria in this setting. The difference between the two examples is therefore that while in the first the benefit in BT is due to the effectiveness of the IGP carried out, in the second the benefit results from the lack of information sharing under AT.

Related Work

Cooperation and coordination is a prevalent theme in multi-agent literature (Kraus 1997). Various mechanisms have been proposed over the years for nurturing and improving cooperation between agents, mostly in settings where the agents are fully cooperative or adhere to the same goal (Rosenfeld et al. 2008; Stone et al. 2010; Gunn and Anderson 2013; Dutta, Jennings, and Moreau 2005). Yet when the agents are fully cooperative, there is no need for solutions of the type discussed in this paper, as the agents will follow any plan that maximizes the social welfare (or any other goal externally set for the group) anyhow. Some works have considered cooperation between self-interested agents that can potentially have conflicting goals (Conitzer 2012; Xie et al. 2010; Wooldridge et al. 2013), in particular aiming at designing norms to guide and constrain agent behavior in order to facilitate cooperation (Haynes, Miles, and Luck 2013; Testerink, Dastani, and Meyer 2013; Hexmoor, Venkata, and Hayes 2006). Other work that considers cooperation between self-interested agents can be found in coalition formation literature (Shehory and Kraus 1998), yet the focus there is mostly on the way coalitions are formed and the division of coalition payoffs, rather than ways of enhancing beneficial individual efforts for the benefit of the coalition.

Group-based cooperation of self-interested agents can also be found in social good allocation games (Aumann 1998; McKelvey and Palfrey 1992; Halonen-Akatwijuka 2012; Schmitz 2015). Here, the use of a threshold-based criteria for increasing social welfare is common (Marks and Croson 1999; Cadsby et al. 2008). Still, the threshold in these works is usually defined over the overall production of the group (i.e., only if the sum of individual contributions is greater than the pre-specified threshold is the group's task is executed) rather than individually as in our case. More importantly, a below-threshold criteria is not an option, as it is doomed to perform poorly in this domain.

The underlying cooperative information gathering model that is used in this paper is rooted in economic search theory where optimal stopping rules are studied for settings where individuals need to gather information on an applicable opportunity when information gathering is costly (Rochlin, Sarne, and Laifenfeld 2012; Grosfeld-Nir, Sarne, and Spiegler 2009; Smith 2011; Lippman and McCall 1976; Alkoby, Sarne, and Das 2015; Rochlin, Sarne, and Zussman 2011). Within this framework, several cooperative information gathering models have been studied, primarily in the context of fully cooperative agents that attempt to maximize the overall utility (Rochlin and Sarne 2013; Gatti 1999; Manisterski, Sarne, and Kraus 2008; Burdett and Malueg 1981; Carlson and McAfee 1984; Rochlin et al. 2016). Criteria setting in these models is therefore irrelevant. Cooperative information gathering by self-interested agents has been studied in several works (Rochlin, Sarne, and Mash 2014; Hazon et al. 2013). These, however, do not suggest any methods for improving the cooperative information gathering in such settings. In our prior work (Rochlin and Sarne 2015) we have introduced a criteria-based mechanism for the cooperative information gathering model focusing on a BT criteria, to which we compare our AT criteria, yielding the result that neither generally dominates the other (and the somehow counter intuitive examples where enabling those who contributed least to benefit from the fruits of cooperation performs better than benefiting those who contributed most).

Discussion and Conclusions

The results reported in this paper provide important evidence for the benefit one may potentially gain from setting criteria for agents to become eligible for benefiting from their cooperation with others. While the criteria prevents some of the agents from benefiting from the cooperative effort (sometimes even without being their fault, as it is possible that an agent will invest all its resources in attempting to comply with the criteria and still fail), its positive effect, if properly set, on individual efforts can have a much greater impact.

The proposed above-threshold cooperative information gathering mechanism encapsulates many benefits compared to the below-threshold one, as discussed throughout the paper: It is more intuitive, it has a closed form solution, it results in lower variance and it is considered to be the most normative choice of the two. Another advantage of the above-threshold mechanism is that it does not requires any enforcement for agents' reporting. With the below-threshold mechanism agents have an incentive to report lower findings, in order to be eligible to benefit from the findings of others as part of the ISP. With the above-threshold mechanism an agent will always have an incentive to report the finding that is most beneficial for the group.

The analysis provided demonstrates that the choice of the criteria to be used should be made carefully, as sometimes it is best to use a below-threshold criteria. This is somehow counter-intuitive, for the reasons given throughout the paper.

In future work we plan to identify and analyze other domains and applications in which better results can be achieved by setting criteria to nurture group-beneficial behaviors through what might seem to be benefiting those offering a somewhat lesser contribution than others.

Acknowledgments

This research was partially supported by the ISRAEL SCIENCE FOUNDATION (grants No. 1083/13) and the ISF-NSFC joint research program (grant No. 2240/15).

References

A. van Wissen, A.; Gal, Y.; Kamphorst, B.; and Dignum, M. 2012. Human–agent teamwork in dynamic environments. *Computers in Human Behavior* 28(1):23–33.

Alkoby, S.; Sarne, D.; and Das, S. 2015. Strategic free information disclosure for search-based information platforms. In *Proceedings of AAMAS'15*, 635–643.

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

Barbulescu, L.; Rubinstein, Z. B.; Smith, S. F.; and Zimmerman, T. L. 2010. Distributed coordination of mobile agent teams: the advantage of planning ahead. In *Proceedings of AAMAS'10*, 1331–1338.

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

Cadsby, C. B.; Croson, R.; Marks, M.; and Maynes, E. 2008. Step return versus net reward in the voluntary provision of a threshold public good: An adversarial collaboration. *Public Choice* 135(3):277–289.

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

Conitzer, V. 2012. Computing game-theoretic solutions and applications to security. In *Proceedings of AAAI'12*, 2106–2112.

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

Dias, M. B., and Sandholm, T. 2004. *TraderBots: A New Paradigm for Robust and Efficient Multirobot Coordination in Dynamic Environments*. Ph.D. Dissertation, Robotics Institute, Carnegie Mellon University.

Dutta, P., and Sen, S. 2003. Forming stable partnerships. *Cognitive Systems Research* 4(3):211–221.

Dutta, P. S.; Jennings, N. R.; and Moreau, L. 2005. Cooperative information sharing to improve distributed learning in multi-agent systems. *Journal of Artificial Intelligence Research* 24:407–463.

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

Grosfeld-Nir, A.; Sarne, D.; and Spiegler, I. 2009. Modeling the search for the least costly opportunity. *EJOR* 197(2):667–674.

Gunn, T., and Anderson, J. 2013. Dynamic heterogeneous team formation for robotic urban search and rescue. *Procedia Computer Science* 19:22 – 31.

Hajaj, C.; Hazon, N.; and Sarne, D. 2015. Improving comparison shopping agents competence through selective price disclosure. *Electronic Commerce Research and Applications* 14(6):563–581.

Halonen-Akatwijuka, M. 2012. Nature of human capital, technology and ownership of public goods. *Journal of Public Economics* 96(1112):939 – 945.

Haynes, C.; Miles, S.; and Luck, M. 2013. Monitoring the impact of norms upon organisational performance: A simulation approach. In *Proceedings of COIN'13*, 103–119.

Hazon, N.; Aumann, Y.; Kraus, S.; and Sarne, D. 2013. Physical search problems with probabilistic knowledge. *Artificial Intelligence* 196:26–52.

Hexmoor, H.; Venkata, S. G.; and Hayes, D. 2006. Modelling social norms in multiagent systems. *Journal of Experimental and Theoretical Artificial Intelligence* 18(1):49–71.

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

Kraus, S.; Shehory, O.; and Taase, G. 2003. Coalition formation with uncertain heterogeneous information. In *Proceedings of AAMAS'03*, 1–8.

Kraus, S. 1997. Negotiation and cooperation in multi-agent environments. *Artificial intelligence* 94(1):79–97.

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

Manisterski, E.; Sarne, D.; and Kraus, S. 2008. Enhancing cooperative search with concurrent interactions. *Journal of Artificial Intelligence Research* 32(1):1–36.

Marks, M. B., and Croson, R. T. 1999. The effect of incomplete information in a threshold public goods experiment. *Public Choice* 99(1):103–118.

McKelvey, R., and Palfrey, T. 1992. An experimental study of the centipede game. *Econometrica* 60(4):803–836.

McMillan, J., and Rothschild, M. 1994. Search. In *Proceedings of Handbook of Game Theory with Economic Applications*, 905–927.

Peled, N.; Bitan, M.; Keshet, J.; and Kraus, S. 2013. Predicting human strategic decisions using facial expressions. In *Proceedings of IJCAI'13*, 2035–2041.

Rochlin, I., and Sarne, D. 2013. Information sharing under costly communication in joint exploration. In *Proceedings of AAAI'13*, 847–853.

Rochlin, I., and Sarne, D. 2015. Constraining information sharing to improve cooperative information gathering. *Journal of Artificial Intelligence Research* 54(1):437–469.

Rochlin, I.; Aumann, Y.; Sarne, D.; and Golosman, L. 2016. Efficiency and fairness in team search with self-interested agents. *Autonomous Agents and Multi-Agent Systems* 30(3):526–552.

Rochlin, I.; Sarne, D.; and Laifenfeld, M. 2012. Coordinated exploration with a shared goal in costly environments. In *Proceedings of ECAI'12*, 690–695.

Rochlin, I.; Sarne, D.; and Mash, M. 2014. Joint search with self-interested agents and the failure of cooperation enhancers. *Artificial Intelligence* 214:45–65.

Rochlin, I.; Sarne, D.; and Zussman, G. 2011. Sequential multilateral search for a common goal. In *Proceedings of IAT'11*, 349–356.

Rosenfeld, A.; Kaminka, G.; Kraus, S.; and Shehory, O. 2008. A study of mechanisms for improving robotic group performance. *Artificial Intelligence* 172(6-7):633–655.

Schmitz, P. W. 2015. Government versus private ownership of public goods: The role of bargaining frictions. *Journal of Public Economics* 132:23 – 31.

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

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

Stone, P.; Kaminka, G. A.; Kraus, S.; and Rosenschein, J. S. 2010. Ad hoc autonomous agent teams: Collaboration without pre-coordination. In *Proceedings of AAAI'10*, 1504–1509.

Tang, Z.; Smith, M. D.; and Montgomery, A. 2010. The impact of shopbot use on prices and price dispersion: Evidence from online book retailing. *International Journal of Industrial Organization* 28(6):579–590.

Testerink, B.; Dastani, M.; and Meyer, J. C. 2013. Norms in distributed organizations. In *Proceedings of COIN'13*, 120–135.

Waldeck, R. 2008. Search and price competition. *Journal of Economic Behavior and Organization* 66(2):347–357.

Wooldridge, M.; Endriss, U.; Kraus, S.; and Lang, J. 2013. Incentive engineering for boolean games. *Artificial Intelligence* 195:418–439.

Xie, S.; Liu, Y.; Zhang, Y.; and Yu, R. 2010. A parallel cooperative spectrum sensing in cognitive radio networks. *IEEE Transactions on Vehicular Technology* 59(8):4079–4092.

Yamamoto, J., and Sycara, K. 2001. A stable and efficient buyer coalition formation scheme for e-marketplaces. In *Proceedings of AGENTS'01*, 576–583.