

The Advantages of Compromising in Coalition Formation with Incomplete Information*

Sarit Kraus

Dept. of CS, Bar-Ilan U.
Ramat-Gan, 52900 Israel
U. Maryland Col. Park, MD
sarit@cs.biu.ac.il

Onn Shehory

IBM Research Lab in Haifa
Haifa U. Mt. Carmel, Haifa
31905 ISRAEL
onn@il.ibm.com

Gilad Taase

Dept. of CS
Bar-Ilan Univ.
Ramat-Gan,
52900 Israel

Abstract

This paper presents protocols and strategies for coalition formation with incomplete information under time constraints. It focuses on strategies for coalition members to distribute revenues amongst themselves. Such strategies should preferably be stable, lead to a fair distribution, and maximize the social welfare of the agents. These properties are only partially supported by existing coalition formation mechanisms. In particular, stability and the maximization of social welfare are supported only in the case of complete information, and only at a high computational complexity. Recent studies on coalition formation with incomplete and uncertain information address revenue distribution in a naïve manner. In this study we specifically refer to environments with limited computational resources and incomplete information. We propose a variety of strategies for revenue distribution, including the strategy in which the agents attempt to distribute the estimated net value of a coalition equally. A variation of the equal distribution strategy in which agents compromise and agree to a payoff lower than their estimated equal share, was specifically examined. Our experimental results show that, under time constraints, the compromise strategy is stable and increases the social welfare compared to non-compromise strategies.

1. Introduction

Coalitions serve as a means for multi-agent collaboration. Agents within coalitions can perform tasks that they might otherwise be unable to perform. Recognizing this, several studies have suggested mechanisms for agent coalition formation, e.g., [7][14][17].

In this paper we consider situations where self-interested agents may benefit from forming coalition when they have incomplete information about each other and there are time constraints. This is the case in the Request For Proposal

(RFP) domain, where some requester business agent issues an RFP—a complex task comprised of sub-tasks—and several service provider agents need to join together to address this RFP. In such environments the RFP value may be common knowledge, however the costs an agent incurs for performing a specific sub-task are unknown to other agents. Additionally, time for addressing RFPs is limited.

These constraints make it difficult to apply traditional coalition formation mechanisms, since those assume complete information, and time constraints are of lesser significance there. Furthermore, coalition formation consists of two main tasks: (i) an agent needs to decide whom to form coalition with, and (ii) the members of a coalition need to agree on the distribution of coalition gains among themselves. Most classical coalition methods address only the distribution task [5]. However, these two tasks are interdependent and both require a combinatorial search. A simultaneous solution of both increases the complexity further, especially in the case of incomplete information. In [8] the first task was assumed for the RFP domain, however the revenue distribution method provided there suggested an equal distribution of profits among coalition members by a trusted agent. Such an approach seems too naïve for self-interested agents and requires a lot of effort on the part of the trusted agent.

In this study we propose several strategies for revenue distribution and integrate them with the coalition selection strategies identified in [8]. Note that due to the incomplete information the agents do not know the actual net profit of a coalition and the distribution strategy can refer only to the estimated profits. In particular, we suggest strategies for revenue distribution where: (1) the estimated net value of the coalition is divided equally among its members; (2) the estimated net value is distributed proportionally to the relative contribution of each member to the value of the coalition; (3) the estimated net value is distributed based on a variation of the Kernel solution concept [1]. We also consider variations of these strategies where coalition members compromise some of their profit computed

* This research was supported in part by NSF grant #IIS0208608 .

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

AAMAS'04, July 19-23, 2004, New York, New York, USA.
Copyright 2004 ACM 1-58113-864-4/04/0007...\$5.00

according to the Kernel, proportional, or equal distribution strategies. These strategies were tested for stability. We also measured the social welfare of the agents (i.e., the sum of their profits) when using these strategies, since stable strategies that yield a high social welfare may encourage self-interested agents to join. Our findings indicate that the strategy in which the agents attempt to distribute the estimated net value of a coalition equally, but each agent is willing to compromise on a payoff lower than its estimated equal share, is stable. Furthermore, this strategy yields the highest social welfare among all the considered strategies.

2. The problem

Given a set of tasks, $\mathcal{T} = \{T_1, \dots, T_n\}$, each task $T_i \in \mathcal{T}$ consists of sub-tasks $T_{i1}, \dots, T_{i r_i}$ and a set of self-interested agents, $\mathcal{A} = \{A_1, \dots, A_m\}$. Each agent is capable of performing only a subset of the subtasks of each task. This partiality is expressed by a boolean function, ϕ . $\phi(A_j, T_{ik})$ evaluates to *true* if A_j is capable of performing T_{ik} , and to *false* otherwise. We assume that ϕ is common knowledge. An agent A_j incurs a cost, b'_{ik} , for executing subtask T_{ik} . Agent $A_{i'}$, $i' \neq i$ does not know b'_{ik} . That is, costs are private knowledge. Agents may however be able to estimate costs of others.

The general problem we study is, given the above, to allow agents to perform tasks and maximize their profits. Since each agent cannot perform a whole task by itself, cooperative task performance is required. This can be achieved by forming coalitions. A coalition \mathcal{C}_T for a task T is a tuple $\langle C, alloc, U \rangle$ where C is a set of member agents, $alloc$ is an allocation function that associates with each subtask of T a member of C such that $alloc(T_{ik}) = A_j$ only if $\phi(A_j, T_{ik}) = \text{true}$. $U = \langle u_1, \dots, u_{|C|} \rangle$ is a payoff distribution vector – u_i is the payoff of A_i . The gross value of a coalition is $V = \sum u_i$. Note that the agents do not know the net value of a coalition (because the costs agents incur are private information), and thus can agree only on the distribution of the payoff V . In our solution, we assume that agents are rational and join coalitions only when they believe this will increase their benefits. For simplicity, we assume that an agent can participate in a single coalition at a time and can perform only one sub-task at a time.

3. Solution approach

To allow agents to form coalitions given the special settings presented above, we have devised a mechanism that consists of a protocol and a set of strategies. To avoid the exponential search for optimal strategies, our solution uses heuristic strategies. Participating agents must adhere to the protocol, and this adherence is enforceable. The use of the suggested strategies in conjunction with the protocol is not

enforced, but we show that the strategies are stable and hence it is reasonable to assume that agents will use them instead of searching for others. The details of the protocol and the strategies follow.

3.1 Coalition formation protocol

The coalition formation protocol, presented initially in [8], is a special type of an auction with an extension for coalition formation. The protocol consists of a central manager and multiple agents that can join in. The manager supports two roles – an auctioneer role and a coalition negotiation manager role. Both roles are neutral trusted third parties that neither discriminate among participating agents nor disclose their private information to others.

The auctioneer publishes the available tasks, collects proposals of coalitions addressing these tasks, determines the winning coalition for each task, and discounts the price of the task over time. The auctioneer also distributes payments to coalition members after they complete task execution, however the decision upon the partition of payments depends on the agreement the members of the coalition arrived at when they formed the coalition.

The auction protocol is used to allocate the tasks in \mathcal{T} to coalitions. The auction is performed in rounds, r_1, r_2, \dots and it ends when there are no more tasks to auction or no more agents that participate in the auction, or the values of the remaining tasks are all nullified by the iterative discount. At the beginning of the auction, the auctioneer announces, for each task $T_i \in \mathcal{T}$, the price, $P(T_i)$, that will be paid to a coalition that will perform T_i and that had formed in the first round of the auction. In each round, the prices of unallocated tasks are reduced by a factor δ (δ is announced by the auctioneer at the beginning of the auction). At each round r , for each unallocated task $T_i \in \mathcal{T}$, there may be zero or more coalitions that propose to perform the task. The auctioneer awards each task to the first competent coalition. Selection among multiple simultaneous proposals is done randomly. The winning coalition is paid $P(T_i)$ upon completing the task. $P(T_i)$ should then be distributed among the members of the coalition. This distribution is the focus of this paper. Note that partial fulfillment of a task yields no payment. Additionally, the submission of a proposal to the auctioneer is binding.

Prior to submitting auction proposals, the agents need to form coalitions. This is performed via negotiation. During the negotiation, agents send and receive proposals for coalitions to be formed. A proposal by an agent A_j specifies a coalition $\mathcal{C}_T = \langle C, alloc, U \rangle$ where $A_j \in C$. The coalition formation negotiation is performed via the negotiation manager. At each auction round the protocol allows only one negotiation round. This is enforced by the manager. At each negotiation round, the manager orders the agents randomly, and the agents perform negotiation actions in that

order. Each agent, in its turn, can either send a proposal for forming a coalition C to all of its members or receive such a formation proposal made to it by another. An agent has only one turn in each round. All offers are valid for one round and thus an agent making an offer must wait until it hears from all of the agents to which it proposed. It cannot accept any other proposal in this round. Note that proposals are all sent via the auctioneer and recorded there. As a result, an attempt to bypass the protocol – sending proposals externally and agreeing on coalitions to be formed – will be detected by the manager and can be penalized for. If all the members of a proposed coalition accept the proposal, the coalition is proposed to the auctioneer for performing a specific task. If a coalition is awarded the task then the members of the coalition quit the negotiation. The agents that have not joined a coalition in a given round continue to negotiate in the next round.

3.2 Ranking strategies

An agent that participates in the protocol presented above needs means to decide which coalitions to propose to which other agents and what revenue distribution to offer. As stated earlier, computing the best strategy to handle such proposals is exponentially complex. To decrease the complexity of the strategies, we propose a two stages decision procedure. First, the agent ranks the possible coalitions. Second, for the top coalition, it computes the revenue distribution. It will offer this coalition to others, and accept only proposals in which its net benefit is at least as high as its net benefit from its computed coalition. We propose strategies for both stages and seek combined strategies that are stable and maximize the social welfare.

Computing and ranking coalitions is performed using information regarding available tasks, their sub-tasks, the payment $P(T_i)$ for the task, the capabilities of other agents and their estimated costs for performing the subtasks. An agent computes the coalitions in which it can jointly address available tasks and then ranks these. Computing coalitions, even without ranking, is exponentially complex, however a cap on coalition size reduces this complexity to polynomial. Fortunately, such a cap is acceptable in many real markets, including the RFP domain. Otherwise, the search itself would require simplifying heuristics to provide feasibility.

In previous research, we have shown that ranking strategies are strongly affected by agents' knowledge. In particular, an agent's knowledge regarding the actual costs other agents incur for executing task plays a major role in this respect. When actual costs of others are not known, agents may estimate these and base their decisions on estimations. In many real markets, some rough estimation based on common knowledge of task costs is available.

Following, we describe two ranking strategies. Elaborate experiments with these have shown good results compared

to a centrally computed optimum and to other strategies. For the completeness of this paper, some of those results are briefly presented here, although they already appear in [8]. The first strategy ranks coalitions by their marginal profit. We thus denote this strategy as *marginal*. To rank a coalition using this strategy, an agent computes the marginal profit of each candidate coalition, and sorts these values. A higher value is ranked higher. The marginal profit of a coalition is the difference between the sum of the (estimated) costs of its members and the value of the task the coalition should perform. This difference is the agent's estimated net value of the task. The intuition of this strategy is that coalitions that have a greater net value may provide a higher utility to the member agents, hence should be preferred over others.

The second strategy ranks coalitions with respect to the expertise of the ranking agent within these coalitions. We denote this strategy as *expert*. An agent is considered an expert with respect to a given coalition if it can perform within sub-tasks that only a few, or none, other members can perform. The *expert* strategy ranks coalitions in which an agent is an expert higher than other coalitions.

Using the proposed strategies, each agent at its turn inspects and ranks possible coalitions. An agent that received proposals compares its share from the best proposal received to its share from its top-rank coalition. If the best proposal is acceptable, it acknowledges acceptance. In case it received no acceptable proposal, or no proposal at all, at the current negotiation round, it sends the coalition with the highest rank as a proposal to the members of the candidate coalition.

3.3 Strategies for revenue distribution

When proposing a coalition $\mathcal{C}_T = \langle C, alloc, U \rangle$, an agent uses the coalition ranking methods presented above to determine the coalition members C and their allocation to subtasks, $alloc$. To compute its proposed U , the agent needs methods for computing revenue distributions. Such methods are presented below.

3.3.1 Equal Distribution. The equal distribution strategy, denoted as *equal*, attempts to offer an equal share of the profit (which is the net revenue) to all agents participating in a coalition. When an agent does not know the exact costs other agents incur for performing their subtasks, it can only offer an estimated equal share. According to *equal*, after selecting the coalition with the highest rank, an agent estimates the sum of costs for all agents participating in the proposed coalition, $B'_C = \sum_j b'_j$, where b'_j is the estimated cost of A_j , summing over all A_j members of C . It then estimates the coalition net revenue, which for performing task T_i in round r is $N' = P(T_i) \delta^r - B'_C$. *Equal* guides the agent to distribute N' evenly among the participating

agents. Each agent will be paid its estimated cost b'_j plus an equal share of N' . That is, A_j 's share will be $b'_j + N'/|C|$.

The *equal* strategy is a simple allocation method, expressing a naive type of fairness. In spite of its simplicity, *equal* has arrived at desirable results in terms of both social welfare and stability, as we later show.

3.3.2 Proportional Distribution. *Equal* offers the same estimated profit to each member of a coalition. In real-world situations, such a revenue allocation might not be realistic. The proportional distribution method, denoted as *proportional*, offers each agent a payoff proportional to its cost of executing its subtask. Intuitively, this reflects the estimated investment needed to perform the subtask, and sometimes also the risks incurred. After selecting a coalition, the agent computes its estimated net benefit N' , as in *equal*. The share of an agent A_j with an estimated cost b'_j will be $b'_j + \frac{N' \cdot b'_j}{\sum b'_i}$, b'_i is the cost A_i incurs performing its subtask in the coalition.

3.3.3 Kernel Distribution. A desired property of a revenue distribution is stability. We consider two notions of stability. The first concept, discussed above, refers to the stability of coalition formation strategies. The second concept of stability refers to formed coalitions, and requires that once a coalition is formed, it will not be worthwhile for a group of agents to break it and form another coalition. Our protocol dictates that once a coalition is formed, it cannot be decomposed. However, this is a restriction that one might prefer to relax. Therefore, a method for computing revenue distributions should address this second notion of coalition stability as well. For evaluating this stability notion, we use the Kernel [1], a well-known game-theoretical stability concept. We also suggest a strategy for revenue distribution based on the Kernel, denoted as *kernel*.

The Kernel is non-empty for any coalition configuration. To find a revenue distribution in the Kernel, we use Stearns' method, which converges to a Kernel point from any given revenue distribution [15]. The details follow:

(1) *Select the best coalition*: rank all coalitions and choose the coalition \mathcal{C}_T with the highest rank (use a ranking strategy); (2) *Create coalition structure S*: assign agents not in \mathcal{C}_T to coalitions to form a structure S such that the sum of the gross values of the coalitions in S is maximized (use hill-climbing); (3) *Initialize payoff vector*: $\forall \mathcal{C} \in S$, distribute $V_{\mathcal{C}}$ arbitrarily among \mathcal{C} members; (4) *Compute the kernel*: use Stearns scheme as follows; (4.1) *Compute the demand function*: $\forall A_i, A_j \in \mathcal{C}_T$, compute the demand function d_{ij} ; (4.1.1) \forall possible coalition R that includes A_i but excludes A_j , compute the excess. The excess $e(R, U)$ of a coalition R is the difference between the net value of R and V_R , the sum of payoffs as suggested by the payoffs in U . The maximum surplus S_{ij} of A_i over A_j is the maximal excess of all possible

coalitions that include A_i but exclude A_j :

$$S_{ij} = \max_{R | A_i \in R, A_j \notin R} e(R, U)$$

(4.1.2) Compute the demand function d_{ij} as follows:

$$d_{ij} = \begin{cases} \min[(s_{ij} - s_{ji})/2, u_j], & \text{if } s_{ij} > s_{ji} \\ 0, & \text{otherwise} \end{cases}$$

(4.2) *Reduce differences*: find agents $A_i, A_j \in \mathcal{C}_T$ with the maximal d_{ij} . If d_{ij} is higher than a certain threshold (set here to 0.1), stop. Otherwise, subtract d_{ij} from A_j 's payoff and add d_{ij} to A_i 's payoff, and then repeat step (4). At the end of this process, U will be in the Kernel.

As seen above, Stearns scheme refers to all possible coalitions and hence requires estimates of the expected profits of all agents, not only those participating in a proposed coalition. It may be difficult to estimate profits of agents that have never formed coalitions. Thus, our Kernel revenue distribution strategy suggests the following: after ranking all possible coalitions and selecting the preferred one, an agent will compute a (near) optimal allocation of other tasks to the remaining agents. Then, the agent will use Stearns' scheme to compute a revenue distribution for the selected coalition. Because of the estimated profits, the near-optimal allocation, and the threshold (in step 4), this is an approximation for a Kernel point.

The demand function computed as part of Stearns' scheme would also be used to compare the stability of the proposed distribution methods. In this respect, a lower average demand function of a method provides evidence that it is more stable.

3.3.4 Compromise. Humans sometimes choose to compromise their righteous profit, if they may benefit from this compromise. We hypothesize that compromise may prove beneficial in the case of coalition revenue distribution. This hypothesis will be checked in our experiments. Compromise can be applied to any of the proposed distribution methods. For instance, *equal* provides agent A_j with a payoff of $b'_j + N'/|C|$. The agent might be satisfied with a part α of its profit, that is, $b'_j + \alpha N'/|C|$, allowing the distribution of $(1 - \alpha)N'/|C|$ among other agents.

Our experiments attempt to find α values that yield high overall utility, and are stable, such that an agent that deviates and requests a payoff using a different α does not increase its profit. Further discussion of the way in which α is determined appears in Section 4.2.1. We evaluated experimentally the proposed strategies for selecting coalitions, accepting proposals and distributing the revenue, as we detail in the following section.

4. Experimental evaluation

In a previous study [8] our goal was to examine the negotiation protocol and compare the strategies suggested

to be used in conjunction with the protocol. We assumed that the benefits will be distributed by the auctioneer. In this study our major goal is to examine strategies for revenue distribution to be used given the protocol and the coalition ranking strategies. Via a series of experiments, we measure the expected gains resulting from the use of the proposed distribution strategies and the stability of each. We consider a strategy to be stable if, given that all agents use it, it will not be worthwhile for an agent to deviate and use another strategy. This notion of stability was tested via experiments in which all agents but one use one strategy and the one agent uses another strategy. This type of stability is referred to as an experimental equilibrium.² The experiments and their results are presented and analyzed below.

4.1 Settings

In our study, experimental settings vary over the number of agents, the number of tasks and the number of subtasks per task. We denote such a setting by a tuple $\langle a, t, s \rangle$, where a, t, s refer to number of agents, number of tasks and number of subtasks, respectively. Given an experimental setting, a specific configuration further requires determining the following parameters: the value of each task, the capabilities of each agent (i.e., which sub-tasks it can perform) and the cost of a given agent to perform each sub-task. The experimental settings we consider are the following: (i) $\langle 6, 2, 3 \rangle$; (ii) $\langle 10, 2, 5 \rangle$; (iii) $\langle 10, 5, 5 \rangle$; (iv) $\langle 12, 3, 4 \rangle$; (v) $\langle 16, 4, 4 \rangle$; (vi) $\langle 16, 5, 4 \rangle$. These provide a variety of agents/tasks combinations. In settings i,ii,iv, and v, the number of agents is equal to the number of possible subtasks to be performed. Thus, the ability to perform all of the tasks depends on the agents' capabilities. In settings iii and vi, the number of subtasks to be performed is larger than the number of agents. Such settings in our model (according to which each agent performs only one subtask at a time), result in some tasks not being performed.

In each experiment, for each of the settings, we randomly generated up to 1,500 configurations. Half of the subtasks within a configuration were "specialized tasks" and the other half were "regular tasks". An agent had a 0.4 probability of being able to perform a regular subtask, but only a 0.15 probability for a specialized one. Task and subtask values and costs were determined as follows. For each subtask, we have randomly selected a mean cost mc with a uniform distribution between 1 and 99. The value of a task is the sum of the means of its subtasks times 1.5, providing an average profit of 50%. The actual cost of a given subtask was randomly drawn from a normal distribution with the mean cost mc , and a deviation of 2.

² An experimental equilibrium, as defined in [3], refers to an equilibrium which is measured experimentally with respect to a given set of strategies. When new strategies are introduced, the equilibrium can be re-computed.

The discount factor δ was set to 0.01. Throughout the experiments, we distinguish between two setting types, referring to the information available to the agents. In the first setting type, referred to here as the "complete information" case, each agent knows the costs of the other agents. In the second setting type, referred to as the "incomplete information" case, the agents know only the mean values of the costs. The capabilities of the agents are common knowledge in both cases.

4.2 Results

The metric we use for evaluating the social welfare of the agents obtained when using various revenue distribution strategies is the ratio between the sum of agent payoffs within a system where agents implement our proposed strategies, and the equivalent centrally computed near optimal sum. Combinatorial complexity prohibits the computation of the optimal cumulative payoff, however a near-optimal value was computed using hill-climbing.

Recall that in a previous study [8] we have examined a rather restrictive revenue distribution method, according to which profits are distributed equally among coalition members, and this distribution is performed by a central manager. In this study, and in particular in the set of experiments we examine alternative revenue distribution strategies. Previous results indicate that, in the case of incomplete information, *expert* was better than *marginal*; in the case of complete information, *marginal* was better. Both strategies performed significantly better with complete information, though *marginal* gained from complete information much more than *expert* did. The following experiments may use these results as a reference.

4.2.1 Compromising. As mentioned above compromising may be beneficial, especially in environments where fast contract closing is important, as in our case. Let $\alpha \in [0..1]$ quantify the willingness of an agent to compromise its profit, $\alpha=1$ refers to no compromise. Smaller α values promote acceptance of other agents' offers, leading to an increase in overall utility. In the extreme case ($\alpha=0$) where each agent demands a zero profit for itself and divides the whole profit between the other agents, all negotiated contracts will be signed immediately, and the average gain of all agents will be high. However, this extreme strategy is unstable, as deviating free riders will profit. We sought α values such that, if all agents follow a strategy that dictates demanding the same α share, it will not be worthwhile for any of them to deviate. We were aware that such α may not exist, nevertheless we did find such α , as discussed below.

In the first experiment of this section, we tested various values of α , from 0 to 1.02. Most of the agents demanded their exact complete share, without compromising, but in each environment, one agent deviated and compromised for

α of its share. We calculated the average utility of all agents, versus the utility of the one deviating agent. We used the strategies according to their superiority: *expert* for incomplete information and *marginal* for complete information. We computed an average of 1500 runs.




Figure 1. Compromise is beneficial

Figure 1 shows the ratio between the utility of the deviating agent and the average utility of all other agents. Clearly, a small compromise increases the profit of the compromising agent. The peak for incomplete information is at $\alpha=0.8$. For the complete information the peak is at $\alpha=0.9$, but the difference between the profits for $\alpha=0.8$ and $\alpha=0.9$ is not significant. When an agent demands more than its rightful share, as for $\alpha=1.02$, it inflicts a sharp decrease on its profits. Note that although compromise increases the gains for both complete and incomplete information, the gains for incomplete information are more significant. This is also apparent in Figures 2 and 3.

Based on the α results above, we studied the marginal and expert strategies subject to compromise where $\alpha=0.8$. We compared two homogenous environments, one where all agents implement $\alpha=0.8$, and one where all agents implement $\alpha=1$. In each, we experimented with both marginal and expert, with incomplete and complete information. We expected the overall utility for $\alpha=0.8$ to be larger than for $\alpha=1$.




Figure 2. Incomplete info. Figure 3. Complete info.

As Figure 2 shows, with incomplete information, compromising for 0.8 of the expected payoff led to a significant increase in agents' average gain (t-test, $p<0.001$). Interestingly, the compromise nullified the advantage of *expert* over *marginal* in the incomplete information case. When the agents compromise, contracts

are signed faster than when they do not. This reduces the effectiveness of *expert*'s ability of resolving conflicts, and therefore *marginal* yielded gains similar to *expert*'s gains. In the complete information case (Figure 3), compromising did not increase *marginal*'s gains. It did increase the gains of *expert*, however these were still lower than those of *marginal*.

Experiments not presented here because of space limitation show that deviating from $\alpha=1$ to $\alpha=0.8$ is significantly profitable (t-test, $p<0.001$). Similar experiments with several α values have shown, in both incomplete and complete information cases, that deviating from $\alpha=0.8$ was not beneficial, hence the $\alpha=0.8$ compromise is a dominant strategy. Furthermore, when using compromise distribution, *marginal* becomes the stable strategy for the case of incomplete information too.

4.2.2 Proportional revenue distribution. We studied the effect of using *proportional*, compared to *equal*. We hypothesized that *equal* would give a higher overall utility, because it will better match the agents' expectations of their share.

The results of experiments not presented here show that the overall profit is slightly higher when using *equal* with compromise ($\alpha=0.8$) than when using *proportional*. In additional experiments deviation was examined. The results (Figure 4) show that deviation from *equal* to *proportional* was not worthwhile (t-test, incomplete info.: $p=0.004$, complete info.: $p=0.04$). However, deviating from *proportional* to *equal* was beneficial (t-test, incomplete: $p<0.001$, complete: $p=0.009$); in that context, *proportional* is not stable, while *equal* is.




Figure 4: Equal/Propo.




Figure 5: Equal/Kernel




Figure 6: Deviation kernel/equal, equal/kernel

4.2.3 Kernel allocation and stability testing. Above, we have examined stability by allowing some agents to deviate

from the majority strategy, checking whether they gained or lost by deviating. This method of stability testing only compares a selected strategy with some other selected ones. However, there might be other, not considered strategy, which performs better than the examined ones. To overcome this, a game-theoretic method for computing equilibrium — the Kernel — is used. In the following experiments, we compare *kernel* to *equal*, focusing on the overall utility and the average demand function.

In the first set of experiments, we compared *equal* and *kernel*; in both cases *marginal* was used for coalition selection. We conducted the experiments for both $\alpha=1$ and $\alpha=0.8$ (compromise). Here, compromise affected both *equal* and *kernel*. The Kernel tends to divide revenues unevenly; very often, a weak agent may be offered a zero payoff. Because with *equal* the agents' expectations match the offers that they get better than with *kernel*, we expected that more contracts will be signed and that the overall profits will be higher with *equal*.

The results (for compromise, seen in Figure 5) confirm our hypothesis. In most cases, the *kernel* average payoff was substantially lesser than the *equal* payoff (t-test, $p<0.001$). Stability in the coalition decomposition sense, measured by means of average demand, was expected to be better (i.e., have lower values) for *kernel* than it would for *equal*. We have computed the average demand function of some homogenous environments, and the results are listed in Table 1. As expected, the *kernel* was found more stable than both *equal* and *proportional*, as its average demand was significantly lower than the average demand of the others. We can also observe that *equal* is more stable than *proportional*, and that environments where all agent compromise are the least stable.

Table 1. Average demand: *kernel* the most stable

Information:	Incomplete	Complete
Equal	4.34	4.86
Proportional	5.33	5.47
Kernel	1.77	1.10
Equal 0.8 Compromise	5.12	5.11
Kernel 0.8 Compromise	2.14	1.76
Equal with Adaptation	4.46	4.84
Proportional with Adaptation	5.29	5.44

We then tested the revenue distribution strategies in heterogeneous environments, where agents deviate from the majority method. We used the dominant $\alpha=0.8$, and studied deviation from *equal* to *kernel* and vice versa. Similar experiments were conducted with no compromise, yet the results were comparable.

Figure 6-right shows that deviating from *equal* to *kernel* is not profitable (t-test, $p<0.001$). This result was not surprising. However, the results of the experiment depicted in Figure 6-left might require some explanations: we see that even if all agents are using *kernel*, it is worthwhile for one agent to deviate and use *equal* (t-test, $p<0.001$). This seems unexpected, since the kernel aims at creating a stable state; but it is not, given the notion of stability that we used.

The *kernel* ensures stability of coalitions that were formed,

such that breaking a coalition will not be profitable; with our protocol, this is unnecessary, as the protocol enforces the stability of formed coalitions. The kernel method might have an important role if the protocol would allow coalitions decomposition.

5. Related Work

Game theory provides various stability concepts for determining distribution of coalition values (see, e.g., [5]) but usually does not consider situations where there is incomplete information about coalition values. There are only few attempts to generalize the stability concepts of coalition formation, such as the core, for situations of asymmetric information [16],[8]. Additionally, game theory does not provide algorithms that agents can use to form coalitions and to reach an agreement on value distribution. Thus, given a specific negotiation protocol for coalition formation, a game theoretic stability concept does not necessarily provide stable strategies. We demonstrated this problem by implementing a revenue distribution strategy based on the Kernel [1], for which we have shown via experiments, that deviation to equal distribution with compromise is beneficial.

Many group formation algorithms for cooperative environments were suggested (e.g., [14],[4]). The revenue division is not important in such settings. Agents may have different views on the environment and tasks, and thus need to compromise in reaching a coalition formation agreement. However, since they try to maximize social welfare, they simply follow system-imposed strategies, with no attempt to deviate. In [17] the problem of coalition formation is addressed for self-interested agents, but in superadditive environments. In [13], solutions were proposed for non-superadditive environments where the value distribution is based on the Kernel. However the value of each coalition is known and the stability of the overall strategies was not considered. Sandholm and Lesser [12] present a coalition formation model for bounded-rational agents and a general classification of coalition games. As in [12] we also allow for varying coalitional values, however we provide the agents with strategies that could be computed in polynomial time. We assume that time is costly, and that agents take the coalition formation time into consideration when deciding on whether to join a coalition. We focus on the stability of the proposed strategies, while Sandholm and Lesser focus on the agents' social welfare. Sandholm et al. [11] discuss the problem of identifying coalition structure that maximizes the sum of the values of coalitions. They neither discuss coalition value distribution nor the stability of forming such coalitions. Griffiths and Luck [2] introduced the notion of clans, a group of agents that share similar objectives, and treat each other favorably when making decision about cooperation. They described mechanisms to

form, maintain and dissolve clans of self-interested agents. However, they do not address deviation either from the clan, or from commitments in the context of the clan. We explicitly address deviation.

All the works we mentioned assume complete information: each of the agents knows the exact value of each possible coalition. For the problem we solve, this assumption does not hold. In real world situations, rarely do other agents know each agent's exact value and costs of fulfilling each task [6]. Therefore, solutions presented in the studies discussed above are inapplicable for our problem. In particular, the methods used to check the stability of a given state require that all agents hold the same beliefs about the state. More related to our work is research on fuzzy and stochastic co-operative games [10]. In such games agents face situations of uncertainty, including, for example, vagueness of expected coalition values and corresponding payoffs. This preliminary research attempts to find formal models to address these problems, while we provide experimental results that present the advantages of using our proposed protocol and strategies.

6. Conclusion

In this paper we consider the problem of coalition formation for cases where groups of self-interested agents can only perform tasks cooperatively. In particular, we consider situations where a complex task is comprised of sub-tasks, and each sub-task should be performed by a different agent. We studied situations in which the costs that an agent incurs for performing a specific sub-task may be unknown to other agents, and time for addressing a task is limited. We focused on the problem of revenue distribution, seeking stable strategies for this distribution. We also aimed at maximizing the social welfare, i.e., the sum of agent payoffs. We found out, via extensive simulation experiments, that the strategy according to which the agents attempt to distribute the estimated net value of a coalition equally, but each agrees to compromise 20% of its estimated equal share, is stable. Further, that strategy yields the highest social welfare when compared to the other strategies that we examined. The stability exhibited by the compromising strategy is surprising, as it contrasts the well-known free-rider phenomenon. The advantages of compromising in fully cooperative multi-agent systems, where agents attempt to maximize the system overall performance, are widely known. However, our results concern self-interested agents who may deviate to increase their own expected rewards. Our hypothesis is that compromise is stable and beneficial under pressing time constraints. In our environment, these constraints are manifested in two ways: firstly, in each time period the value of a given possible coalition is discounted; secondly, tasks are awarded to coalitions based on the first-formed-

first-awarded criterion. These factors encourage the formation of coalitions as early as possible. Apparently, agents that are not willing to compromise are not able to join a beneficial coalition. Future work should study additional multi-agent environments where compromising of self-interested agents is stable and beneficial.

7. References

- [1] Davis, M., and Maschler, M. The Kernel of a cooperative game. *Naval Research Logistics Quart.*, 12 (1965), 223-259.
- [2] Griffiths, N. and Luck, M. Coalition formation through motivation and trust, *Proc. of AAMAS 2003*, 17-24.
- [3] Goldman, C., Kraus, S., and Shehory, O. On Experimental Equilibria Strategies for Selecting Sellers and Satisfying Buyers, *Decision Support Systems Journal*. 2003.
- [4] Jennings, N. R. Controlling cooperative problem solving in industrial multi-agent systems using joint intentions. *AI Journal*, 75(2) (1995), 1-46.
- [5] Kahan, J. P., and Rapoport A. *Theories of coalition formation*. Lawrence Erlbaum Associates, Hillsdale NJ, 1984.
- [6] Klusch, M. and Gerber, A. Issues of dynamic coalition formation among agents. *Proc. KSCO-2002*, France.
- [7] Klusch, M., and Shehory, O. A polynomial kernel-oriented coalition formation algorithm for rational information agents. *Proc. ICMAS 1996*, 157-164.
- [8] Kraus, S., Shehory, O., and Taase, G. Coalition formation with uncertain heterogeneous information. *Proc. AAMAS'03 (Melbourne Australia, 2003)*, 1-8.
- [9] Kwasnica, A. Bayesian Implementable of Efficient and Core Allocations. *Penn. State U. Working Paper*, 2002.
- [10] Mares, M. Fuzzy coalition structures, *Fuzzy Sets and Systems*, 114 (2000), 23-33.
- [11] Sandholm, T., Larson, K., Andersson, M., Shehory, O. and Tohme F. Coalition structure generation with worst case guarantees, *AI Journal*, 111(1-2) (1999), 209-238.
- [12] Sandholm, T., and Lesser, V. Coalition formation among bounded rational agents, *AI Journal* 94(1-2, (1997), 99-137,
- [13] Shehory, O., and Kraus, S. Feasible formation of coalitions among autonomous agents, *Computational Intelligence*, 15(3) (1999), 218-251.
- [14] Shehory, O., and Kraus, S. Methods for task allocation via agent coalition formation, *AI Journal*, 101 (1-2) (1998), 165-200.
- [15] Sterns, R. E. Convergent transfer schemes for n-person games. *Transactions of the American Math Society*, 134, (1968), 449-459.
- [16] Vohra, R. Incomplete Information, Incentive Compatibility and the Core, *J. of Economic Theory*, 86 (1999), 123-147.
- [17] Zlotkin, G., and Rosenschein, J. Coalition, cryptography, and stability: mechanisms for coalition formation in task oriented domains, *Proc. AAAI '94 (Seattle WA, 1994)*, 432-437.