

Axiomatic Characterization of Task Oriented Negotiation*

Dongmo Zhang

Intelligent Systems Laboratory
University of Western Sydney, Australia
d.zhang@uws.edu.au

Abstract

This paper presents an axiomatic analysis of negotiation problems within task-oriented domains (TOD). We start by applying three classical bargaining solutions of Nash, Kalai-Smorodinsky and Egalitarian to the domains of problems with a pre-process of randomization on possible agreements. We find out that these three solutions coincide within any TOD and can be characterized by the same set of axioms, which specify a solution of task oriented negotiation as an outcome of dual-process of maximizing cost reduction and minimizing workload imbalance. This axiomatic characterization is then used to produce an approximate solution to the domain of problems without randomization on possible agreements.

1 Introduction

Negotiation or bargaining is a typical form of interaction between intelligent agents. Initiated by Nash's seminal work [Nash, 1950], the study of negotiation in game theory has reached high sophistication with a variety of models and solutions. The theory has been extensively applied to economics, management science, social science as well as computer science [Binmore *et al.*, 1992; Thomson, 1994; Rosenschein and Zlotkin, 1994; Jennings *et al.*, 2001].

Game-theoretic account of bargaining consists of two typical models: *axiomatic model* (cooperative theory) and *strategic model* (non-cooperative theory). The axiomatic model specifies a bargaining problem as a one-shot game with complete information and characterizes bargaining solutions axiomatically [Thomson, 1994]. The strategic model devices explicit construction of negotiation procedures and identifies bargaining outcomes as equilibria [Binmore *et al.*, 1992]. The attempt to establish the relationship between the two models is known as the *Nash program*.

The axiomatic model of negotiation is domain-dependent, which means that a solution can have totally different characteristics with different domain of problems. Given the Nash

solution as an example, the axiomatic systems for convex domain, comprehensive domain, finite domain, and ordinal domain are significantly different (see Section 6 for a brief summary). This feature causes the axiomatic theory of bargaining rather complicated and the research exceedingly interesting.

Task Oriented Domain (TOD), introduced by Zlotkin and Rosenschein (1993), represents a class of negotiation problems that can be abstracted as task sharing among autonomous agents. More specifically, the domain specifies a certain kind of interaction among agents in which the agents have sets of tasks to carry out, and can exchange tasks and share in their execution. A range of real-world applications, such as parcel delivery, database queries, job allocation, can be described in this domain [Rosenschein and Zlotkin, 1994].

As originated from computer science and AI, the task oriented domain has specific features that differentiate it from those that have been studied in game theory. The domain is not convex even after randomization, therefore the traditional axiomatizations for the Nash solution, the Kalai-Smorodinsky (KS) solution and the Egalitarian solution are not applicable [Nash, 1950; Kalai and Smorodinsky, 1975]. The modern extensions of these solutions to the non-convex domain are inapplicable as well because the domain is not comprehensive [Conley and Wilkie, 1991; 1996; Xu and Yoshihara, 2006]. Those solutions specially designed for discrete domains are also not suitable for task oriented negotiation [Mariotti, 1998; Özgür Kibris and Sertel, 2007]. As we will see, the above mentioned three solutions coincide within TODs and their axiomatic characterization is significantly different from the ones for any other domains.

The paper is organized as the following. Section 2 recalls the basic concepts on TOD from [Rosenschein and Zlotkin, 1994]. Section 3 presents a cooperative model of task oriented negotiation. In Section 4, we redefine and characterize three classical bargaining solutions in mixed deals. In Section 5, we demonstrate how the axiomatic characterization can be used to produce an approximate solution in pure deals. In the final two sections, we conclude the paper with a discussion of related work.

2 Task oriented domains

A task oriented domain specifies a negotiation situation in which a group of agents have to decide how to cooperate for

*This research was partially supported by the Australian Research Council through Discovery Project DP0988750.

reducing execution costs. More formally, we have the following definition.

Definition 1 [Zlotkin and Rosenschein, 1993] A *task oriented domain* (TOD) is a tuple $\langle \mathcal{T}, \mathcal{A}, c \rangle$ where:

1. \mathcal{T} is a finite set of possible tasks;
2. $\mathcal{A} = \{A_1, A_2\}$ is the set of agents;
3. $c : 2^{\mathcal{T}} \rightarrow \mathbb{R}^+$ is a monotonic function that maps a set of possible tasks to a non-negative real number, satisfying $X \subseteq Y \subseteq \mathcal{T}$ implies $c(X) \leq c(Y)$;
4. $c(\emptyset) = 0$.

In other words, a task oriented domain specifies a set of possible tasks, a group of agents who are capable of carrying out any possible combination of tasks, and a cost function which determines the cost of execution of each set of tasks. Different from the original setting in [Zlotkin and Rosenschein, 1993], we restrict the set of possible tasks to be finite and the number of agents to be two.

Rosenschein and Zlotkin (1994) have demonstrated that a number of real-world negotiation scenarios can be specified in task oriented domains. All these scenarios satisfy the following common assumption:

Subadditivity: $c(X \cup Y) \leq c(X) + c(Y)$ for any $X, Y \in 2^{\mathcal{T}}$.

In other words, regrouping tasks could reduce execution costs. This motivates the agents to negotiate for a better task distribution. Throughout this paper, we assume that all the domains we consider are subadditive.

3 Task oriented negotiation

In order to facilitate analysis of task oriented negotiation, Rosenschein and Zlotkin (1994) has introduced a few concepts to model the problems. Let us recall the basic concepts that will be used in this paper.

Given a TOD $\langle \mathcal{T}, \mathcal{A}, c \rangle$, a pair (T_1, T_2) is an *encounter* within the TOD if, for each $i \in \{1, 2\}$, T_i is a subset of \mathcal{T} . We denote the set of all encounters within the TOD by $\Sigma(\mathcal{T}, \mathcal{A}, c)$.

Let $T = (T_1, T_2)$ be an encounter, a pair (D_1, D_2) is a *pure deal* of T if $D_1 \cup D_2 = T_1 \cup T_2$. In other words, a pure deal is a redistribution of tasks in T among the agents. We let $\Omega(T)$ be the set of all pure deals of T . Note that T itself is always a pure deal of T , which is referred to as the *conflict deal* of the encounter. Also, $(T_1 \cup T_2, \emptyset)$ and $(\emptyset, T_1 \cup T_2)$ are pure deals of T .

Let (D_1, D_2) be a pure deal of T . Any expression in the form of “ $(D_1, D_2) : p$ ” is called a *mixed deal* of T if p is a probability, i.e., $(0 \leq p \leq 1)$. Intuitively, $(D_1, D_2) : p$ represents a probabilistic redistribution of the tasks in (D_1, D_2) in the way that agent 1 is assigned D_1 with probability p and D_2 with $1 - p$ while agent 2 receives the sets of tasks with the inverse probability. The set of all the mixed deals of T is denoted by $\overline{\Omega}(T)$. In certain context, we view a pure deal (D_1, D_2) as the mixed deal $(D_1, D_2) : 1$. Moreover, the mixed deal $(T_1 \cup T_2, \emptyset) : p$ is called to be an *all-or-nothing* deal of T . For more intuitive discussions on the above concepts, see [Rosenschein and Zlotkin, 1994] Chapters 3 & 4.

We would like to remark that the operation of randomization to generate a mixed deal is significantly different from the way that Nash uses to convexify a feasible set, where an “anticipation” is generated from *two* anticipations in the form: $pD' + (1 - p)D''$. As a result, any feasible set becomes convex after Nash’s randomization. However, the set of mixed deals is not necessarily convex (see Figure 1).

Next we extend the cost function of a TOD to mixed deals. For each encounter T within a TOD $\langle \mathcal{T}, \mathcal{A}, c \rangle$, we define a cost function $C : \overline{\Omega}(T) \rightarrow \mathbb{R}^2$ as follows: for each $D = (D_1, D_2) : p \in \overline{\Omega}(T)$, C defines the cost of the mixed deal to each agent, i.e., $C(D) = (C_1(D), C_2(D))$, where

$$C_1(D) = p * c(D_1) + (1 - p) * c(D_2)$$

$$C_2(D) = (1 - p) * c(D_1) + p * c(D_2)$$

Note that for a mixed deal $(D_1, D_2) : p$, we always have $C_1(D) + C_2(D) = c(D_1) + c(D_2)$. In other words, the process of randomizing a pure deal redistributes cost (workload) between two agents but does not reduce cost. To reduce cost, we need to redistribute tasks. In fact, there are two major factors that determine the outcome of a task oriented negotiation: *cost reduction* and *workload distribution*. We will use these two factors as the criteria to compare deals.

For any two mixed deals D and D' of an encounter, we say that D *dominates* D' , denoted by $D \succeq D'$, if and only if $C_i(D) \leq C_i(D')$ for all $i \in \{1, 2\}$. We say that D *strongly dominates* D' , denoted by $D \succ D'$, if and only if $D \succeq D'$ and there is an i such that $C_i(D) < C_i(D')$. Moreover, we say that D is *equivalent* to D' , written as $D \approx D'$ if and only if $D \succeq D'$ and $D' \succeq D$. Obviously, \approx is an equivalence relation on $\overline{\Omega}(T)$. Based on the concept of domination, we define the Pareto set of $\overline{\Omega}(T)$ as:

$$\overline{P}(T) = \{D \in \overline{\Omega}(T) : \text{there is no } D' \in \overline{\Omega}(T) \text{ s.t. } D' \succ D\} \quad (1)$$

We can also compare two deals in terms of workload distribution. We consider a deal to be “*fairer*” than the other if its workload distribution between agents is closer to the original allocation of the encounter. Formally, for any $D', D'' \in \overline{\Omega}(T)$, we write $D' \triangleright_T D''$ if and only if $\text{dist}(D', T) < \text{dist}(D'', T)$, where

$$\text{dist}(D, T) = |(C_1(D) - C_2(D)) - (c(T_1) - c(T_2))| \quad (2)$$

Intuitively, $D' \triangleright_T D''$ means that the workload distribution of D' is closer to T than that of D'' . Note that we use the difference of costs to each agent to measure the workload distribution between the agents.

All the above concepts can also be defined on pure deals. To save space, we omit their definitions.

The following lemmas describe the properties of Pareto set, which will be frequently used in the proof of other theorems. The first lemma is a theorem in [Rosenschein and Zlotkin, 1994] (Theorem 9), which says that Pareto optimal deals maximize cost reduction. The second lemma asserts that any Pareto optimal deal has an all-or-nothing representative. We omit their proof here.

Lemma 1 [Rosenschein and Zlotkin, 1994] *For any $(D_1, D_2) : p \in \overline{P}(T)$, $c(D_1) + c(D_2) = c(D_1 \cup D_2)$.*

Lemma 2 $(D_1, D_2) : p \in \overline{P}(T)$ if and only if there exists an all-or-nothing deal D' such that $D \approx D'$.

4 Negotiation solutions in mixed deals

In this section, we consider the solutions that take mixed deals as output. We apply three most outstanding bargaining solutions in game theory: the *Nash solution*, the *Kalai-Smorodinsky (KS) solution* and the *Egalitarian solution*¹, to TODs and characterize these solutions.

4.1 Game-theoretic bargaining model

Negotiation analysis aims to provide a clear-cut prediction of the outcomes of any negotiation situation within the domain under consideration. In game theory, a bargaining solution is a function that assigns to each bargaining game a single agreement or a set of agreements depending on the underlying domains.

Formally, a *bargaining game* is a pair (S, d) , where $S \subseteq \mathbb{R}^2$ is the feasible set that can be derived from possible agreements and $d \in S$ stands for the disagreement point. A *bargaining solution* f is a function that assigns to each bargaining game (S, d) a unique point of S or a subset of S .

Given an encounter T within a TOD $\langle \mathcal{T}, \mathcal{A}, c \rangle$, we can define a game-theoretic bargaining game by using Rosenschein and Zlotkin's utility function: for each $D \in \overline{\Omega}(T)$, $u_1(D) = c(T_1) - C_1(D)$ and $u_2(D) = c(T_2) - C_2(D)$. We let $S = \{(u_1(D), u_2(D)) : (D_1, D_2) \in \overline{\Omega}(T)\}$ and $d = (0, 0)$. Then (S, d) is the bargaining game that corresponds to the encounter T . One may think that the task oriented negotiation problem could have been solved if we apply a game-theoretic solution to each of the bargaining games. Unfortunately this is not the case. As we have mentioned in the introduction, all game-theoretic negotiation solutions are domain-dependent. As we will see in Example 1 and 2, all three solutions we consider in this section give different results in mixed deals and in pure deals.




Figure 1: An example of bargaining game in task oriented domain with mixed deals

Figure 1 shows an example of bargaining game in task oriented domain with mixed deals (the dots represent pure deals

¹see [Thomson, 1994] for a comprehensive introduction of these solutions.

and the dashed lines represent mixed deals. We omit the original example due to space limitation). It is easy to see that such a bargaining game is neither convex nor comprehensive, which means that the existing characterizations for the classical bargaining solutions are not applicable to this domain.

4.2 Three classical solutions

Although the axiomatic characterizations for classical game-theoretic solutions are not applicable to task oriented negotiation, the solutions themselves can be applied to the domain of problems even though their behavior will be very likely different. In this subsection, we redefine the three classical bargaining solutions in the context of task oriented domains.

Definition 2 Given a TOD $\langle \mathcal{T}, \mathcal{A}, c \rangle$, a *negotiation solution f in mixed deals* is a function that assigns to each encounter $T \in \Sigma(\mathcal{T}, \mathcal{A}, c)$ a set of mixed deals, i.e., $f(T) \subseteq \overline{\Omega}(T)$.

Note that a negotiation solution is set-valued, which means that the solution outputs a set of mixed deals as its prediction. This is because there might be several equivalent ways to redistribute tasks.

Now we define a few specific solutions. To facilitate these definitions, let $\overline{I}(T)$ be the individual rational deals of T , i.e.,

$$\overline{I}(T) = \{D \in \overline{\Omega}(T) : c(T_i) - C_i(D) \geq 0 \text{ for } i = 1, 2\} \quad (3)$$

First, we define the Nash solution in the task oriented domain [Nash, 1950] (also see [Rosenschein and Zlotkin, 1994] p.50).

Definition 3 A negotiation solution f in mixed deals is the *Nash solution* if for all $T \in \Sigma(\mathcal{T}, \mathcal{A}, c)$,

$$f(T) = \arg \max_{D \in \overline{I}(T)} (c(T_1) - C_1(D))(c(T_2) - C_2(D)) \quad (4)$$

where $\arg \max$ means the arguments of \max function, i.e., the maximizers of the product of the utilities of two agents.

Next, we apply the Egalitarian solution to our domain:

Definition 4 A negotiation solution f in mixed deals is the *Egalitarian solution* if for all $T \in \Sigma(\mathcal{T}, \mathcal{A}, c)$,

$$f(T) = \arg \max_{D \in \overline{I}(T)} \{v : c(T_1) - C_1(D) = c(T_2) - C_2(D) = v\} \quad (5)$$

that is, the deals that maximize and divide equally the utilities of two agents.

Finally, we simulate Kalai-Smorodinsky's solution [Kalai and Smorodinsky, 1975]. To this end, let $(a_1, a_2) = (\max_{D \in \overline{I}(T)} (c(T_1) - C_1(D)), \max_{D \in \overline{I}(T)} (c(T_2) - C_2(D)))$, i.e., the ideal point. Then we have

Definition 5 A negotiation solution f in mixed deals is the *Kalai-Smorodinsky solution* if for all $T \in \Sigma(\mathcal{T}, \mathcal{A}, c)$,

$$f(T) = \arg \max_{D \in \overline{I}(T)} \{v : \begin{aligned} (c(T_1) - C_1(D)) &= a_1 v \text{ and} \\ (c(T_2) - C_2(D)) &= a_2 v \end{aligned} \} \quad (6)$$

In other words, the solution gives the individual rational deals that maximize the utilities on the diagonal from $(0, 0)$, the utilities of the conflict deal, to the ideal point (a_1, a_2) .

Example 1 Let $\langle T, \mathcal{A}, c \rangle$ be a TOD where $\mathcal{T} = \{a, b, c\}$, $\mathcal{A} = \{A_1, A_2\}$ and the cost function is the following:

$$\begin{aligned} c(\emptyset) &= 0, c(\{a\}) = 1, c(\{b\}) = 2, c(\{c\}) = 3, \\ c(\{a, b\}) &= 2, c(\{a, c\}) = 4, c(\{b, c\}) = 4, \\ c(\{a, b, c\}) &= 5. \end{aligned}$$

Consider an encounter $T = (\{a, b\}, \{b, c\})$ within the TOD. It is not hard to know that the Nash solution, KS solution and Egalitarian solution give the same set of mixed deals with the elements that are all equivalent to the all-or-nothing deal $(\{a, b, c\}, \emptyset) : 0.3$. \square

Lemma 3 For any encounter T within a TOD, $f^N(T) = f^{KS}(T) = f^E(T)$, where f^N, f^{KS}, f^E are the Nash solution, the KS solution and the Egalitarian solution, respectively.

Lemma 4 Let f be any of the three solutions defined above. For any encounter T within a TOD and any $D, D' \in f(T)$, $D \approx D'$. In other words, each of the three solutions gives unique outcome modulo \approx .

The above results reflect the special features of task oriented domains. The first result says that these three solutions are actually the same. This is not surprising. Since a task oriented domain assumes that all agents share the same cost function, it is impossible for an agent to vary its utility to gain more negotiation power. In other words, all the agents have the same negotiation power. Therefore they can share the available utility equally (note that this does not mean that all the agents will be assigned the same amount of tasks or the tasks with the same cost). On the contrary, the second result is a bit surprising because typical solutions on non-convex domains are multiple-valued [Mariotti, 1998]. We have such a property because of the assumption of subadditivity on the cost function and randomization on pure deals.

4.3 Characterization

We have seen above that the task oriented domain is a specific domain that unifies three classical bargaining solutions. This indicates that the characterization of these solutions within a TOD requires specific axioms. It is well known that the Egalitarian solution can be characterized in convex domain by the axioms: *Pareto optimality, scale invariance, symmetry* and *strong monotonicity* [Thomson, 1994]. We notice that scale invariance is not applicable to a TOD because the agents share the same cost function. Without this axiom, symmetry becomes useless. In addition, strong monotonicity is also hard to apply because an arbitrary subset of the set of mixed deals does not necessarily correspond to another encounter. However, it is easy to see that Pareto optimality and individual rationality (implied by other axioms in game theory) are true with any TOD. In short, the following axioms will characterize all three solutions within a TOD.

IR: $D \in f(T)$ implies $D \succeq T$. (Individual rationality)

PO: $f(T) \subseteq \overline{P}(T)$. (Pareto optimality)

NV: $f(T) \neq \emptyset$. (Non-vacuity)

Eq: $D \in f(T)$ and $D \approx D'$ imply $D' \in f(T)$. (Equivalence)

WB: $D \in f(T)$ and $D' \triangleright_T D$ imply $D' \notin \overline{\Omega}(T)$. (Workload balance)

The first two axioms are standard requirements for negotiation solutions, which are self-explanatory. The other three are specific to the task oriented negotiation. NV says that the solution guarantees an output (obviously the worst case is the conflict deal). Eq says that all deals are treated the same if their costs to each agent are the same. WB says that a solution should guarantee that the workload distribution of the final agreement must be the closest to the original allocation. Note that this property implies that if an agent is originally allocated heavier tasks than the other (measured in their costs), the agent should receive also heavier tasks in the final agreement.

Theorem 1 Given a TOD $\langle T, \mathcal{A}, c \rangle$, a negotiation function in mixed deals is the Egalitarian solution (therefore, the Nash solution and the KS solution) if and only if it satisfies IR, PO, NV, Eq and WB.

It is easy to see that the key components of the characterization are the two optimality axioms: *PO* and *WB*. *PO* requires a solution to be the best in cost reduction while *WB* requires the solution to be the best to match the original workload distribution.

5 Negotiation solutions in pure deals

In the previous section we allow to randomize a pure deal so that the tasks can be divided in any portion to balance workload. If we disable randomization by restricting a solution to pure deals, the problem will become much harder. To see the problem, let's consider Example 1 again.

Example 2 Consider the negotiation encounter in Example 1. If we restrict the solutions introduced in Section 4.2 to pure deals, only the Nash solution is non-empty, which is $\{(\{a\}, \{b, c\}), (\{b\}, \{a, c\}), (\{a, b\}, \{c\})\}$. In fact, these deals are exactly all the deals that are Pareto optimal and individual rational. The Egalitarian solution and Kalai-Smorodinsky solution are empty. Therefore none of the solutions offers useful prediction.

The reason that causes the failure of all the solutions is that, without mixed deals, we are not always able to divide tasks to each agent to match the original workload distribution due to the indivisibility of atomic tasks. In fact, this is not a problem particularly to task oriented negotiation. Almost every negotiation domain has the same problem: a solution exists in an idealized domain, say convex hell or comprehensive hell, but cannot be found in its original domain. In this section, we demonstrate how to use the characterization we presented in the previous section to construct an approximate solution for the task oriented negotiation problems. First, we redefine negotiation solution by restricting the outcomes to be pure deals.

Given a TOD $\langle T, \mathcal{A}, c \rangle$, a pure deal negotiation solution F is a function that assigns to each encounter $T \in \Sigma(T, \mathcal{A}, c)$ a set of pure deals of T , i.e., $F(T) \subseteq \Omega(T)$.

Definition 6 A pure deal negotiation solution F is the *minimal distance solution* if for all $T \in \Sigma(\mathcal{T}, \mathcal{A}, c)$,

$$F(T) = \arg \min_{D \in NS(T)} \text{dist}(D, T), \quad (7)$$

where $NS(T) = P(T) \cap I(T)$, i.e., the negotiation set of T . $P(T)$, $I(T)$ and $\text{dist}(D, T)$ can be redefined in pure deals by using the equations (1), (3) and (2), respectively.

The following result indicates that the minimal distance solution is an approximation of the Egalitarian solution in mixed deals (we omit the proof of the proposition).

Proposition 1 Let f be the Egalitarian solution in mixed deals on a TOD $\langle \mathcal{T}, \mathcal{A}, c \rangle$. Then for any $T \in \Sigma(\mathcal{T}, \mathcal{A}, c)$,

$$f(T) = \arg \min_{D \in \overline{P}(T) \cap \overline{I}(T)} \text{dist}(D, T)$$

Example 3 Consider the negotiation encounter in Example 1 again. It is easy to calculate that the minimal distance solution of the encounter is:

$$F(T) = \{(\{a\}, \{b, c\}), (\{a, b\}, \{c\})\}. \quad \square$$

Notice that different from the mixed deal solution, $F(T)$ is not unique even after modulo operation. However, the worst case is two modulo equivalence.

Lemma 5 For any encounter T in a TOD, the minimal distance solution $F(T)$ is non-empty and has at most two elements modulo \approx .

The following theorem shows that the minimal distance solution can be characterized by the same set of axioms for the mixed deal solutions.

Theorem 2 Given a TOD $\langle \mathcal{T}, \mathcal{A}, c \rangle$, a negotiation function in pure deals is the minimal distance solution if and only if it satisfies the following properties:

IR: $D \in F(T)$ implies $D \succeq T$. (Individual rationality)

PO: $F(T) \subseteq P(T)$. (Pareto optimality)

NV: $F(T) \neq \emptyset$. (Non-vacuity)

Eq: $D \in f(T)$ and $D \approx D'$ imply $D' \in f(T)$. (Equivalence)

WB: $D \in F(T)$ and $D' \triangleright_T D$ imply $D' \notin P(T)$. (Workload balance)

This theorem again exhibits another specific feature of task oriented domains that different solutions share the same axiomatization. Note that the axioms may play different roles in different characterization. For instance, in Theorem 1, WB forces the solution to be zero distance to T while in Theorem 2, it just requires the solution to be as close to T as possible.

6 Related work

Axiomatic analysis of bargaining situations started from Nash (1950). Nash proposed a bargaining solution and characterized it with a set of axioms. Nash's characterization only applies to convex domains, which can be implemented via randomizing possible agreements. Kaneko (1980) extended Nash's characterization to non-convex domain but enforced continuity on the domain meanwhile allowing the solution to be set-valued. This result was refined by a few

other authors so that the requirement for convexity is replaced by comprehensiveness without need to amend Nash's original axioms [Conley and Wilkie, 1996; Zhou, 1996; Xu and Yoshihara, 2006]. The result was also extended to the Kalai-Smorodinsky solution and Egalitarian solution [Conley and Wilkie, 1991; Hougaard and Tvede, 2003; Xu and Yoshihara, 2006]. However, these results are restricted to continuous domains. Characterizing a bargaining solution on discrete or finite domains is much harder mostly because the idealized points that satisfy certain axioms, say Nash's, may not exist in the feasible set. Among a few authors, Mariotti and Lahiri proposed two different characterizations of the Nash solution in finite domains [Mariotti, 1998; Lahiri, 2003]. Kibris and Sertel characterized a few non-standard solutions for finite ordinal domains [Özgür Kibris and Sertel, 2007]. Unfortunately none of these results can be applied to the task oriented domains. Mariotti's axioms require to operate a feasible set by adding or removing a point. Such an operation is not allowed for our domain. Kibris and Sertel's axioms also apply set operations on feasible sets. Lahiri's characterization applies only on a special finite domain, in which the task oriented problems are not describable.

This work was developed based on Rosenschein and Zlotkin's framework [Zlotkin and Rosenschein, 1993; Rosenschein and Zlotkin, 1994]. The basic concepts of task oriented domains, pure deals and mixed deals are extended from their work. We would like to remark that the special design of randomization for mixed deals is the key to most of the elegant results. Different from Nash's approach, it does not convexify a domain but can guarantee the uniqueness of solution. It maximally keeps the link to original deals so that the original deals are still operable (see Lemma 1) while gains the benefit of divisibility of tasks.

7 Conclusion and discussions

We have presented an axiomatic analysis of negotiation problems within task oriented domains. We found that such a domain has a number of interesting features. Firstly, three classical bargaining solutions coincide when they are applied to a TOD with mixed deals but diverge if their outcomes are restricted to pure deals. Secondly, all these solutions in mixed deals have unique value by equivalence and their pure deal approximation has at most two values. Thirdly, the solutions in mixed deals and their pure deal approximation share the same axiomatic characterization. All these special features originate from the specific setting of task oriented domains. The combination of simplicity of cost function setting and requirement for dual optimality gives the domains rich properties with manageable complexity and makes the research deeply interesting.

Nash's axiomatic model laid on the foundation of bargaining theory. It is also one of the most fundamental models in modern economic theory [Rubinstein, 2000]. An axiomatic model of bargaining fully specify a domain and captures the nature of problems in the domain. It achieves great generality by avoiding any specification of bargaining processes therefore can serve as a guideline for devising variety of specific

negotiation procedures or solutions for specific applications (Section 5 is a simple example of such an application).

Negotiation and bargaining has been a research area in economics and social science for nearly sixty years. The research has been mostly motivated by problems from economics and social science. Problems raised from computer science or AI can be appreciably different from those from other disciplines. Besides task oriented domains, similar domains of problems can also be found in AI literature, such as state oriented domains [Rosenschein and Zlotkin, 1994], argumentation-based negotiation [Kraus *et al.*, 1998], contract negotiation [Dunne, 2005], and etc. Axiomatic analysis on these domains will not only help us to gain a better understanding of these problems but also would make a contribution towards the game-theoretic research of bargaining.

References

- [Binmore *et al.*, 1992] Ken Binmore, Martin J. Osborne, and Ariel Rubinstein. Noncooperative models of bargaining. In R.J. Aumann and S. Hart, editors, *Handbook of Game Theory with Economic Applications*, volume 1, chapter 7, pages 180–225. Elsevier, 1992.
- [Conley and Wilkie, 1991] John P. Conley and Simon Wilkie. The bargaining problem without convexity: extending the Egalitarian and Kalai-Smorodinsky solutions. *Economics Letters*, 36(4):365–369, 1991.
- [Conley and Wilkie, 1996] John P. Conley and Simon Wilkie. An extension of the Nash bargaining solution to nonconvex problems. *Games and Economic Behavior*, 13:26–38, 1996.
- [Dunne, 2005] Paul E. Dunne. Extremal behaviour in multiagent contract negotiation. *Journal of Artificial Intelligence Research*, 23:41–78, 2005.
- [Hougaard and Tvede, 2003] Jens Leth Hougaard and Mich Tvede. Nonconvex n-person bargaining: efficient maxmin solutions. *Economic Theory*, 21:81–95, 2003.
- [Jennings *et al.*, 2001] Nick R. Jennings, Peyman Faratin, Alessio R. Lomuscio, Simon Parsons, Carles Sierra, and Michael Wooldridge. Automated negotiation: prospects, methods and challenges. *International Journal of Group Decision and Negotiation*, 10(2):199–215, 2001.
- [Kalai and Smorodinsky, 1975] Ehud Kalai and Meir Smorodinsky. Other solutions to Nash’s bargaining problem. *Econometrica*, 43(3):513–518, 1975.
- [Kaneko, 1980] Mamoru Kaneko. An extension of the Nash bargaining problem and the Nash social welfare function. *Theory and Decision*, 12:135–148, 1980.
- [Kraus *et al.*, 1998] Sarit Kraus, Katia Sycara, and Amir Evenchik. Reaching agreements through argumentation: a logical model and implementation. *Artificial Intelligence*, 104:1–69, 1998.
- [Lahiri, 2003] Somdeb Lahiri. Axiomatic characterization of the nash and kalai-smorodinsky solutions for discrete bargaining problems. *Pure Mathematics and Applications*, 14(3):207–220, 2003.
- [Mariotti, 1998] Marco Mariotti. Nash bargaining theory when the number of alternatives can be finite. *Social Choice and Welfare*, 15:413–421, 1998.
- [Nash, 1950] John Nash. The bargaining problem. *Econometrica*, 18(2):155–162, 1950.
- [Özgür Kibris and Sertel, 2007] Özgür Kibris and Murat R. Sertel. Bargaining over a finite set of alternatives. *Social Choice and Welfare*, 28:421–437, 2007.
- [Rosenschein and Zlotkin, 1994] Jeffrey S. Rosenschein and Gilad Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*. The MIT Press, 1994.
- [Rubinstein, 2000] Ariel Rubinstein. *Economics and Language: Five Essays*. Cambridge University Press, 2000.
- [Thomson, 1994] William Thomson. Cooperative models of bargaining. In R.J. Aumann and S. Hart, editors, *Handbook of Game Theory*, volume 2, chapter 35, pages 1237–1284. Elsevier, 1994.
- [Xu and Yoshihara, 2006] Yongsheng Xu and Naoki Yoshihara. Alternative characterizations of three bargaining solutions for non-convex problems. *Games and Economic Behavior*, 57(1):86–92, 2006.
- [Zhou, 1996] Lin Zhou. The Nash bargaining theory with non-convex problems. *Econometrica*, 65(3):681–685, 1996.
- [Zlotkin and Rosenschein, 1993] Gilad Zlotkin and Jeffrey S. Rosenschein. A domain theory for task oriented negotiation. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, pages 416–422. Morgan Kaufmann, 1993.

Appendix: Proof of Theorems

Note: Due to space limitation, we omit the proof of a few theorems and outline the proof of the others.

Proof of Lemma 3: All solutions divide available utility equally between the two agents (see [Rosenschein and Zlotkin, 1994] Theorem 11 for the proof of the Nash solution. The proof for the other two is obvious). \square

Proof of Lemma 4: By Lemma 3 we can assume that f is the Nash solution. Firstly, it is not hard to prove that $f(T)$ is non-empty and each deal in $f(T)$ is Pareto optimal. Next, by Lemma 2, any Pareto optimal deal is equivalent to an all-or-nothing deal. Then we can verify that only one all-or-nothing deal can be the maximizer of the product of utilities of two agents. \square

Proof of Theorem 1: “ \Rightarrow ” The proof of IR and Eq is trivial. PO is implied by Lemma 4. WB holds because $dist(D, T) = 0$ for any deal D in the Egalitarian solution.

“ \Leftarrow ” Given any encounter T within a TOD, by NV, there exists a deal D of T such that $D \in f(T)$. PO and Lemma 2 implies that D is equivalent to an all-or-nothing deal D' . By Eq, we have $D' \in f(T)$. Let $D'' = (T_1 \cup T_2, \emptyset) : p$, where $p = \frac{c(T_1 \cup T_2) + c(T_1) - c(T_2)}{2c(T_1 \cup T_2)}$. Since $dist(D'', T) = 0$, WB implies that $dist(D', T) = 0$. It then follows that $c(T_1) - C_1(D') = c(T_2) - C_2(D')$. In addition, D' is individual rational and Pareto optimal. Let $f^E(T)$ be the Egalitarian solution. We then have $D' \in f^E(T)$. By Eq, $D \in f^E(T)$. Therefore $f(T) \subseteq f^E(T)$. However, Lemma 4 says that $f^E(T)$ has only one element modulo \approx . Thus $f(T)$ must be equal to $f^E(T)$. \square

Proof of Lemma 5: Assume that there are three pure deals D^1, D^2, D^3 in $F(T)$. Among these three, there must be at least two deals that make $(c(D_1) - c(D_2)) - (c(T_1) - c(T_2))$ either ≥ 0 or ≤ 0 if replace D with one of them. Without losing generality, we assume that $(c(D_1^1) - c(D_2^1)) - (c(T_1) - c(T_2)) \geq 0$ and $(c(D_1^2) - c(D_2^2)) - (c(T_1) - c(T_2)) \geq 0$. Note that $dist(D^1, T) = dist(D^2, T)$. Therefore $(c(D_1^1) - c(D_2^1)) - (c(T_1) - c(T_2)) = (c(D_1^2) - c(D_2^2)) - (c(T_1) - c(T_2))$. It turns that $c(D_1^1) - c(D_2^1) = c(D_1^2) - c(D_2^2)$. If $c(D_1^1) \neq c(D_1^2)$, say $c(D_1^1) < c(D_1^2)$, we have $c(D_2^1) > c(D_2^2)$ because D^1 and D^2 are Pareto optimal. This implies that $c(D_1^1) - c(D_2^1) < c(D_1^2) - c(D_2^2)$, a contradiction. Therefore $c(D_1^1) = c(D_1^2)$. By Pareto optimality, we have $D^1 \approx D^2$. \square

Proof of Theorem 2: Straightforward from the definition of the minimal distance solution. \square