# Negotiation and Conflict Resolution in Non-Cooperative Domains

**Gilad Zlotkin***
**Jeffrey S. Rosenschein**
Computer Science Department, Hebrew University
Givat Ram, Jerusalem, Israel
gilad@humus.huji.ac.il, jeff@humus.huji.ac.il

## Abstract

In previous work [Zlotkin and Rosenschein, 1989a], we have developed a negotiation protocol and offered some negotiation strategies that are in equilibrium. This negotiation process can be used only when the "negotiation set" (NS) is not empty. Domains in which the negotiation sets are never empty are called cooperative domains; in general *non-cooperative* domains, the negotiation set is sometimes empty.

In this paper, we present a theoretical negotiation model for rational agents in general non-cooperative domains. Necessary and sufficient conditions for cooperation are outlined. By redefining the concept of utility, we are able to enlarge the number of situations that have a cooperative solution. An approach is offered for conflict resolution, and it is shown that even in a conflict situation, partial cooperative steps can be taken by interacting agents (that is, agents in fundamental conflict might still agree to cooperate up to a certain point).

A Unified Negotiation Protocol is developed that can be used in all cases. It is shown that in certain borderline cooperative situations, a partial cooperative agreement (i.e., one that does not achieve all agents' goals) might be preferred by all agents, even though there exists a rational agreement that would achieve all their goals.

## Introduction

The subject of negotiation has been of continuing interest in the distributed artificial intelligence (DAI) community [Smith, 1978; Rosenschein and Genesereth, 1985; Durfee, 1988; Malone *et al.*, 1988; Sycara, 1988; Sycara, 1989; Kuwabara and Lesser, 1989; Conry *et al.*, 1988]. The operation of cooperating, intelligent autonomous agents would be greatly enhanced if they were able to communicate their respective desires and compromise to reach mutually beneficial agreements. The work described in this paper follows the general direction of [Rosenschein and Genesereth, 1985; Zlotkin and Rosenschein, 1989a] in treating negotiation in the spirit of game theory, while altering game theory assumptions that are irrelevant to DAI.

Previous work [Zlotkin and Rosenschein, 1989a] discussed inter-agent negotiation protocols and negotiation strategies that were in equilibrium, but could only be used if the so-called "negotiation set" [Harsanyi, 1977] was not empty. Cooperative domains are those in which NS is never empty; in this paper, we present a theoretical negotiation model for general *non-cooperative* domains (where NS might be empty).

## General Definitions

Two autonomous agents $A$ and $B$ share the same world; this world is in some initial state $s$. Each agent wants the world to satisfy a set of goal conditions.

**Definition 1 Goals**

- *The* goal *of agent $i \in \{A, B\}$, $g_i$, is a set of predicates that agent $i$ wants the world to satisfy.*
- *$G_i$ stands for the set of world states that satisfy all the predicates in $g_i$.*

Both agents have the same set of operations OP that they can perform. An operation $o$ in OP moves the world from one state to another; it is a function $o\colon \mathrm{ST} \to \mathrm{ST}$ where ST is the set of all possible world states.

**Definition 2 Plans**

- *A* one-agent plan *to move the world from state $s$ to state $f$ in* ST *is a list $[o_1, o_2, \ldots, o_n]$ of operations from* OP *such that $f = o_n(o_{n-1}(\ldots o_1(s)\ldots))$.*
- *A* joint plan *to move the world from state $s$ to state $f$ in* ST *is a pair of one-agent plans $(P_A, P_B)$ and a schedule.*

A schedule is a partial order over the union of actions in the two one-agent plans. It specifies that some actions cannot be taken until other actions are completed; because it is a *partial* order, it of course allows

*Supported in part by the Leibniz Center for Research in Computer Science.

simultaneous actions by different agents. If the initial state of the world is $s$ and each agent $i$ executes plan $P_i$ according to the schedule, then the final state of the world will be $f$. We will sometimes write $J$ to stand for a joint plan $(P_A, P_B)$.

**Definition 3 Costs**

- *There exists a* cost function, Cost: OP $\rightarrow \mathbb{N}$.
- *For each one-agent plan* $P = [o_1, o_2, \ldots, o_n]$, Cost($P$) *is defined to be* $\sum_{k=1}^{n} \text{Cost}(o_k)$.
- *For each joint plan* $J = (P_A, P_B)$, $\text{Cost}_i(J)$ *is defined to be* Cost($P_i$).

Note that cost is a function over an operation—it is independent of the state in which the operation is carried out. This definition, however, is not critical to the subsequent discussion. Our theory is insensitive to the precise definition of any single operation's cost. What is important is the ability of an agent to measure the cost of a one-agent plan, and the cost of one agent's part of a joint multi-agent plan.

**Definition 4 Best Plans**

- $s \rightarrow f$ *is the minimal* Cost *one-agent plan that moves the world from state* $s$ *to state* $f$. *If a plan like this does not exist, then* $s \rightarrow f$ *is undefined.*
- $s \rightarrow F$ *(where* $s$ *is a world state and* $F$ *is a* set *of world states) is the minimal* Cost *one-agent plan that moves the world from state* $s$ *to one of the states in* $F$:

$$\text{Cost}(s \rightarrow F) = \min_{f \in F:\, s \rightarrow f \text{ is defined}} \text{Cost}(s \rightarrow f)$$

**Example: The Blocks World Domain**
There is a table and a set of blocks. A block can be on the table or on some other block, and there is no limit to the height of a stack of blocks. However, on the table there are only a bounded number of slots into which blocks can be placed. There are two operations in this world: PickUp($i$) — Pick up the top block in slot $i$ (can be executed whenever slot $i$ is not empty), and PutDown($i$) — Put down the block which is currently being held into slot $i$. An agent can hold no more than one block at a time. Each operation costs 1.[1]

## Underlying Assumptions

In [Zlotkin and Rosenschein, 1989a], we introduced several assumptions that are in force for our discussion here as well (the final two assumptions were implicit in previous work):[2]

1. **Utility Maximizer:** *Each agent wants to maximize his expected utility.*
2. **Complete Knowledge:** *Each agent knows all relevant information.*
3. **No History:** *There is no consideration given by the agents to the past or future; each negotiation stands alone.*
4. **Fixed Goals:** *Though the agents negotiate with one another over operations, their goals remain fixed.*
5. **Bilateral Negotiation:** *In a multi-agent encounter, negotiation is done between a pair of agents at a time.*

## Deals and the Negotiation Set

The agents negotiate on a joint plan that brings the world to a state that satisfies both agents' goals.

**Definition 5 Deals**

- *A* Pure Deal *is a joint plan* $(P_A, P_B)$ *that moves the world from state* $s$ *to a state in* $G_A \cap G_B$.
- *A* Deal *is a mixed joint plan* $(P_A, P_B){:}\,p$; $0 \leq p \leq 1 \in \mathbb{R}$.

*The semantics of a Deal is that the agents will perform the joint plan* $(P_A, P_B)$ *with probability* $p$, *or the symmetric joint plan* $(P_B, P_A)$ *with probability* $1 - p$.

- *If* $\delta = (J{:}\,p)$ *is a* Deal, *then* $\text{Cost}_i(\delta)$ *is defined to be* $p\text{Cost}_i(J) + (1-p)\text{Cost}_j(J)$ *(where* $j$ *is* $i$*'s opponent).*
- *If* $\delta$ *is a* Deal, *then* $\text{Utility}_i(\delta)$ *is defined to be* $\text{Cost}(s \rightarrow G_i) - \text{Cost}_i(\delta)$.

*The utility for an agent from a deal is simply the difference between the cost of achieving his goal alone and his expected part of the deal.*

- *A* Deal $\delta$ *is* individual rational *if, for all* $i$, $\text{Utility}_i(\delta) \geq 0$.
- *A* Deal $\delta$ *is* pareto optimal *if there does not exist another* Deal *which dominates it—there does not exist another* Deal *which is better for one of the agents and not worse for the other.*
- *The* negotiation set NS *is the set of all the deals that are both individual rational and pareto optimal.*

These definitions of an individual rational deal, a pareto optimal deal, and the negotiation set NS are standard definitions from game theory and bargaining theory (see, for example, [Luce and Raiffa, 1957; Nash, 1950; Harsanyi, 1977]).

[1] Note that actions are relative to the current state, and don't name a particular block for manipulation. This emphasizes the role of the schedule in ensuring that the two agents move the correct blocks at the correct times.

[2] See [Zlotkin and Rosenschein, 1990b] for a discussion of how the Fixed Goals assumption can be relaxed. The relaxation of the Complete Knowledge assumption, in specific circumstances, was treated in [Zlotkin and Rosenschein, 1989a]. Future work will further examine the consequences of removing one or more of these assumptions, such as the No History assumption and the Bilateral Negotiation assumption.

## Conditions for Cooperation

A necessary condition for NS to be non-empty is that there is no contradiction between the two agents' goals, i.e., $G_A \cap G_B \neq \emptyset$.[3] This condition is not sufficient, however, because even when there is no contradiction between agents' goals, there may still be a *conflict* between them. In such a conflict situation, any joint plan that satisfies the union of goals will cost one agent (or both) more than he would have spent achieving his own goal in isolation (that is, no deal is individual rational).

**Example:** The initial state can be seen at the left in Figure 1. $g_A$ is "The Black block is at slot 2 but not on the table" and $g_B$ is "The White block is at slot 1 but not on the table".

In order to achieve his goal alone, each agent has to execute one PickUp and then one PutDown; $\text{Cost}(s \rightarrow G_i) = 2$. The two goals do not contradict each other, because there exists a state in the world which satisfies them both, as can be seen on the right side of Figure 1. There does not exist a joint plan that moves the world from the initial state to a state that satisfies the two goals with total cost less than 8—that is, no deal is individual rational.

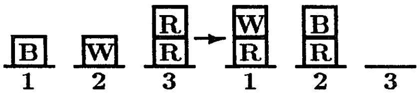


Figure 1: Conflict exists even though union of goals is achievable

The existence of a joint plan that moves the world from its initial state $s$ to a state in $G_A \cap G_B$ is a necessary condition for NS to be non-empty. When this condition is not true, we will call it a *conflict* situation. Ways in which this conflict can be resolved will be discussed in the Conflict Resolution section below.

**Definition 6 Sum and Min Conditions**

- *A joint plan J will be said to satisfy the* sum condition *if*

$$\sum_{i \in \{A,B\}} \text{Cost}(s \rightarrow G_i) \geq \sum_{i \in \{A,B\}} \text{Cost}_i(J).$$

- *A joint plan J will be said to satisfy the* min condition *if*

$$\min_{i \in \{A,B\}} \text{Cost}(s \rightarrow G_i) \geq \min_{i \in \{A,B\}} \text{Cost}_i(J).$$

**Theorem 1** *There exists a joint plan that moves the world from its initial state s to a state in $G_A \cap G_B$ and also satisfies the sum and the min conditions, if and only if* NS $\neq \emptyset$.

[3] All the states that exist in the intersection of the agents' goal sets might, of course, not be reachable given the domain of actions that the agents have at their disposal. See [Zlotkin and Rosenschein, 1989b] for an example of a domain in which such a situation can occur.

**Proof.** For the proof of this theorem and subsequent theorems, see [Zlotkin and Rosenschein, 1990a]. □

When the conditions of Theorem 1 are true, we will say that the situations are cooperative.

## Redefinition of Utility

In non-conflict situations, if neither the min nor the sum conditions are true, then in order for the agents to cooperatively bring the world to a state in $G_A \cap G_B$, at least one of them will have do more than if he were alone in the world and achieved only his own goals. Will either one of them agree to do extra work? It depends on how important each $g_i$ is to agent $i$, i.e., how much $i$ is willing to pay in order to bring the world to a state in $G_i$.

### The Worth of a Goal

**Definition 7** *Let $W_i$ be the maximum* expected *cost that agent i is willing to pay in order to achieve his goal $g_i$.*

We assume that such an upper bound exists. There may be situations and domains in which there is no limit to the cost that an agent is willing to pay in order to achieve his goal—he would be willing to pay *any* cost (see [Zlotkin and Rosenschein, 1989b]). That situation, however, is beyond the scope of this paper.

The declaration of Utility can be usefully altered as follows:

**Definition 8** *If $\delta$ is a deal, then* $\text{Utility}_i(\delta)$ *is defined to be* $W_i - \text{Cost}_i(\delta)$.

The utility for an agent of a deal is the difference between $W_i$ and the cost of his part of the deal. If an agent achieves his goal alone, his utility is the difference between the worth of the goal and the cost that he pays to achieve the goal.

**Theorem 2** *If in Definition 6 we change every occurrence of* $\text{Cost}(s \rightarrow G_i)$ *to $W_i$, then Theorem 1 is still true.*

### Types of Interactions

Before the redefinition of utility, we had two possible situations for agent interaction: *conflict* and *cooperative*. A conflict situation implied a contradiction between the agents' goals, or a cost to achieving the union of their goals that was so high, no deal was individual rational.

Now that utility has been redefined, we have three possible situations for agent interaction: conflict, *compromise*, and cooperative.

- A *conflict* situation is one in which (as before) the negotiation set is empty—no individual rational deals exist.
- A *compromise* situation is one where there are individual rational deals. However, agents would prefer to be alone in the world, and to accomplish their

goals alone. Since they are forced to cope with the presence of other agents, they will agree on a deal. All of the deals in NS are better for both agents than leaving the world in its initial state $s$.

- A *cooperative* situation is one in which there exists a deal in the negotiation set that is preferred by both agents over achieving their goals alone. Here, every agent welcomes the existence of the other agents.

When the negotiation set is not empty, we can distinguish between compromise and cooperative situations using the following criterion. If for all $i$, $W_i \leq \text{Cost}(s \rightarrow G_i)$ and NS $\neq \emptyset$, then it is a cooperative situation; otherwise, it is a compromise situation.[4]

## Conflict Resolution

What can be done when the agents are in a conflict situation?

If we dropped Assumption 3 ("No History"), then we could offer some mechanism in which agents can "buy their freedom" by making a promise to their opponent regarding future actions. In this case, they will negotiate over the price of freedom. A discussion of altering utilities through promises, however, is beyond the scope of this paper.

A simpler solution would be for the agents to flip a coin in order to decide who is going to achieve his goal and who is going to be disappointed. In this case they will negotiate on the probabilities (weightings) of the coin toss. If they run into a conflict during the negotiation (fail to agree on the coin toss weighting), the world will stay in its initial state $s$.[5]

Utility for agent $i$ in general is the difference between the worth for $i$ of the final state of the world and the cost that $i$ spends in order to bring the world to its final state.

If agent $i$ wins the coin toss, then he can reach his goal. In this case, his utility is $W_i$ (the worth of his goal) minus the cost he has to spend in order to bring the word to a state that satisfies his goal. If agent $i$ loses the coin toss, his opponent is going to bring the world to a state that satisfies his opponent's goal. This state will not satisfy $g_i$ (otherwise it would not be a conflict situation). The final state of the world in this case is worth 0 to agent $i$, but he is not going to spend *anything* to bring the world to this state, so his total utility in the case where he loses the coin toss is 0.

If the agents agree to flip a coin with weighting $q$, then the utility for agent $i$ of such a deal is $q_i(W_i - \text{Cost}(s \rightarrow G_i))$, where $q_A = q; q_B = 1 - q$.

**Example:** There is one block at slot 1. $g_A$ is "The block is at slot 2" and $g_B$ is "The block is at slot 3"; $W_A = 12$ , and $W_B = 22$. The agents will agree here on the deal that will give them the same utility—to flip a coin with weighting $\frac{2}{3}$. This deal will give them each a utility of $\frac{20}{3}$.[6]

### Cooperation in Conflict Resolution

The agents may find that, instead of simply flipping a coin in a conflict situation, it is better for them to cooperatively reach a new world state (not satisfying either of their goals) and *then* to flip the coin in order to decide whose goal will ultimately be satisfied.

**Example:** One agent wants the block currently in slot 1 to be in slot 2; the other agent wants it to be in slot 3. In addition, both agents share the goal of swapping the two blocks currently in slot 4 (i.e., reverse the stack's order). See the left side of Figure 2. Assume that $W_A = W_B = 12$. The cost for an agent of achieving his goal alone is 10. If the agents decide to flip a coin in the initial state, they will agree on a weighting of $\frac{1}{2}$, which brings them a utility of 1 (i.e., $\frac{1}{2}(12 - 10)$). If, on the other hand, they decide to do the swap cooperatively (at cost of 2 each), bringing the world to the state shown on the right of Figure 2, and *then* flip a coin, they will still agree on a weighting of $\frac{1}{2}$, which brings them an overall utility of 4 (i.e., $\frac{1}{2}(12 - 2 - 2)$).

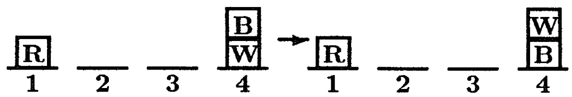


Figure 2: Cooperation up to a certain point

**Definition 9** *A* Semi-Cooperative Deal *is a tuple* $(t, J, q)$ *where* $t$ *is a world state,* $J$ *is a mixed joint plan that moves the world from the initial state* $s$ *to state* $t$, *and* $0 \leq q \leq 1 \in \mathbb{R}$ *is the weighting of the coin toss—the probability that agent* $A$ *will achieve his goal.*

The semantics of such a deal is that the two agents will perform the mixed joint plan $J$, and will bring the world to state $t$; then, in state $t$, they will flip a coin with weighting $q$ in order to decide who continues the plan towards their own goal.

**Definition 10**

$$\begin{aligned}\text{Utility}_i(t, J, q) &= q_i(W_i - \text{Cost}_i(J) - \text{Cost}(t \rightarrow G_i)) \\ &\quad -(1 - q_i)\text{Cost}_i(J) \\ &= q_i(W_i - \text{Cost}_i(t \rightarrow G_i)) - \text{Cost}_i(J)\end{aligned}$$

[4] An example of a compromise situation can be found in Figure 1 when $W_i$ is greater than 4.

[5] There is a special case where the initial state $s$ already satisfies one of the agent's goals, let's say agent $A$ ($s$ cannot satisfy both goals since then we would not have a conflict situation). In this case, the only agreement that can be reached is to leave the world in state $s$. Agent $A$ will not agree to any other deal and will cause the negotiation to fail.

[6] We have $\frac{2}{3}(12 - 2) = \frac{1}{3}(22 - 2) = \frac{20}{3}$.

## Unified Negotiation Protocol (UNP)

In cooperative and compromise situations, the agents negotiate on deals that are mixed joint plans, $J\!:\!p$ (cooperative deals). In a conflict situation, the agents negotiate on deals of the form $(t, J, q)$ (semi-cooperative deals).

We would like to find a Unified Negotiation Protocol (UNP) that the agents can use in any situation. The main benefit would be that the agents would not have to know (or even to agree), prior to the negotiation process, on the type of situation that they are in. Determining whether the situation is cooperative or not may be difficult. An agent may not have full information at the beginning of a negotiation; he may gain more information during the negotiation, for example, from the deals that his opponents are offering, and from computations he himself is doing in order to generate the next offered deal. Agents may only know near the end of a negotiation just what kind of situation they are in.

The semi-cooperative deals $(t, J, q)$ are general enough so that, with some minor changes in the definition of utility, they may be used in the Unified Negotiation Protocol. A *cooperative* deal which is a mixed joint plan $J\!:\!p$ can also be represented as $(J(s), J\!:\!p, 0)$ where $J(s)$ is the final world state resulting from the joint plan $J$ when the initial state is $s$. $J(s)$ is in $G_A \cap G_B$, so the result of the coin flip at state $J(s)$ does not really matter (since none of the agents would want to change the state of the world anyway).

What we advocate is for agents to negotiate always using semi-cooperative deals. A cooperative agreement can still be reached (when the situation is cooperative) because the cooperative deals are a subset of the semi-cooperative deals.

**Definition 11**

- *If $(t, J, q)$ is a semi-cooperative deal, then $f_i$ will be defined as the final state of the world when agent $i$ wins the coin toss in state $t$. $f_i = (t \rightarrow G_i)(t) \in G_i$.*
- *$W_i(f_j) = W_i$ when $f_j \in G_i$, otherwise it is 0.*
- $\text{Utility}_i(t, J, q) = q_i(W_i - \text{Cost}_i(t \rightarrow G_i)) + (1 - q_i)W_i(f_j) - \text{Cost}_i(J)$
- *Two deals $d_1, d_2$ (cooperative or semi-cooperative) will be said to be* equivalent *if* $\forall i$ $\text{Utility}_i(d_1) = \text{Utility}_i(d_2)$. *The calculation of the utility of each deal is done according to the type of the deal (cooperative or semi-cooperative).*

**Theorem 3** *If* $\forall i\ W_i \geq \text{Cost}(s \rightarrow G_i)$, *then* $\text{NS} \neq \emptyset$.

If $W_i < \text{Cost}(s \rightarrow G_i)$ then agent $i$ cannot even achieve his goal alone. This does not necessarily mean that NS is empty—Theorem 3 stated in the opposite direction is not true.

**Theorem 4** *For a semi-cooperative deal* $(t, J, q) \in \text{NS}$, *if there exists an $i$ such that $f_i \in G_A \cap G_B$, then this semi-cooperative deal is equivalent to some cooperative deal.*

It is easy to see that whenever $f_A, f_B \notin G_A \cap G_B$, then the definition of utility in Definition 10 is the same as that in Definition 11.

### UNP in a Cooperative Situation

In a cooperative situation, there is always an individual rational cooperative deal, where both agents' goals are satisfied. One might expect that in such a situation, even if the agents use the Unified Negotiation Protocol, they will agree on a semi-cooperative deal that is equivalent to the cooperative deal, i.e., both goals would be achieved. Surprisingly, this is not the case: there might exist a semi-cooperative deal that dominates all cooperative deals and does *not* achieve both agents' goals. See the example below.

It turns out that this is a borderline situation, brought about because $W_i$ is low. As long as $W_i$ is high enough, any semi-cooperative deal that agents agree on in a cooperative situation will be equivalent to a cooperative deal.

**Example:** The initial situation in Figure 3 consists of 5 duplications of the example from Figure 1, in slots 1 to 15. In addition, two slots (16 and 17) each contain a stack of 2 blocks. $g_A$ is "Black blocks are in slots $2, 5, 8, 11$ and 14 but not on the table; the blocks in slots 16 and 17 are swapped" (i.e., each tower is reversed). $g_B$ is "White blocks are in slots $1, 4, 7, 10$ and 13 but not on the table; the blocks in slots 16 and 17 are swapped".

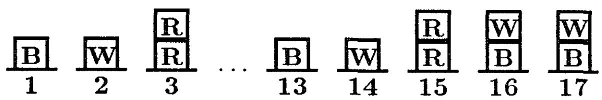


Figure 3: Semi-Cooperative Agreement in a Cooperative Situation

For all $i$, $\text{Cost}(s \rightarrow G_i) = 26 = (2 \times 5) + (8 \times 2)$. Let $J$ be the minimal cost joint plan that achieves both goals. The cooperative deal $J\!:\!\frac{1}{2}$ satisfies the min and the sum conditions, because for all $i$, $\text{Cost}_i(J\!:\!\frac{1}{2}) = 24 = \frac{1}{2}((8 \times 5) + (4 \times 2))$. This situation is cooperative. For all $i$, $\text{Utility}_i(J\!:\!\frac{1}{2}) = 26 - 24 = 2$. Let $t$ be the state where the blocks in slots 16 and 17 are swapped, and the other slots are unchanged. Let $T$ be the minimal cost joint plan that moves the world to state $t$. For all $i$, $\text{Utility}_i(t, T\!:\!\frac{1}{2}, \frac{1}{2}) = \frac{1}{2}(26 - (2 \times 5)) - (2 \times 2) = 4$. The semi-cooperative deal $(t, T\!:\!\frac{1}{2}, \frac{1}{2})$ thus dominates the cooperative deal $J\!:\!\frac{1}{2}$ even though it is a cooperative situation.

## Conclusions

We have presented a theoretical negotiation model that encompasses both cooperative and conflict situations. Necessary and sufficient conditions for cooperation were outlined. By redefining the concept of utility, a new boundary type of interaction, a *compromise* situation, was demarcated. A solution was offered for conflict resolution, and it was shown that even in a conflict situation, partial cooperative steps can be taken by interacting agents. A Unified Negotiation Protocol was developed that can be used in all cases, whether cooperative, compromise, or conflict. It was shown that in certain borderline cooperative situations, a partial cooperative agreement (i.e., one that does not achieve all agents' goals) might be preferred by all agents.

## References


[Conry *et al.*, 1988] Susan E. Conry, Robert A. Meyer, and Victor R. Lesser. Multistage negotiation in distributed planning. In Alan H. Bond and Les Gasser, editors, *Readings in Distributed Artificial Intelligence*, pages 367–384. Morgan Kaufmann Publishers, Inc., San Mateo, California, 1988.

[Durfee, 1988] Edmund H. Durfee. *Coordination of Distributed Problem Solvers.* Kluwer Academic Publishers, Boston, 1988.

[Harsanyi, 1977] John C. Harsanyi. *Rational Behavior and Bargaining Equilibrium in Games and Social Situations.* Cambridge University Press, Cambridge, 1977.

[Kuwabara and Lesser, 1989] Kazuhiro Kuwabara and Victor R. Lesser. Extended protocol for multistage negotiation. In *Proceedings of the Ninth Workshop on Distributed Artificial Intelligence*, pages 129–161, Rosario, Washington, September 1989.

[Luce and Raiffa, 1957] R. Duncan Luce and Howard Raiffa. *Games and Decisions.* John Wiley & Sons, Inc., New York, 1957.

[Malone *et al.*, 1988] Thomas W. Malone, Richard E. Fikes, and M. T. Howard. Enterprise: A market-like task scheduler for distributed computing environments. In B. A. Huberman, editor, *The Ecology of Computation*, pages 177–205. North-Holland Publishing Company, Amsterdam, 1988.

[Nash, 1950] John F. Nash. The bargaining problem. *Econometrica*, 28:155–162, 1950.

[Rosenschein and Genesereth, 1985] Jeffrey S. Rosenschein and Michael R. Genesereth. Deals among rational agents. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pages 91–99, Los Angeles, California, August 1985.

[Smith, 1978] Reid G. Smith. *A Framework for Problem Solving in a Distributed Processing Environment.* PhD thesis, Stanford University, 1978.

[Sycara, 1988] Katia P. Sycara. Resolving goal conflicts via negotiation. In *Proceedings of the Seventh National Conference on Artificial Intelligence*, pages 245–250, St. Paul, Minnesota, August 1988.

[Sycara, 1989] Katia P. Sycara. Argumentation: Planning other agents' plans. In *Proceedings of the The Eleventh International Joint Conference on Artificial Intelligence*, pages 517–523, Detroit, Michigan, August 1989.

[Zlotkin and Rosenschein, 1989a] Gilad Zlotkin and Jeffrey S. Rosenschein. Negotiation and task sharing among autonomous agents in cooperative domains. In *Proceedings of the The Eleventh International Joint Conference on Artificial Intelligence*, pages 912–917, Detroit, Michigan, August 1989. The International Joint Conference on Artificial Intelligence.

[Zlotkin and Rosenschein, 1989b] Gilad Zlotkin and Jeffrey S. Rosenschein. Negotiation and task sharing in a non-cooperative domain. In *Proceedings of the Ninth Workshop on Distributed Artificial Intelligence*, pages 307–327, Rosario, Washington, September 1989.

[Zlotkin and Rosenschein, 1990a] Gilad Zlotkin and Jeffrey S. Rosenschein. Negotiation and conflict resolution in non-cooperative domains. Technical Report 90-6, Computer Science Department, Hebrew University, Jerusalem, Israel, 1990.

[Zlotkin and Rosenschein, 1990b] Gilad Zlotkin and Jeffrey S. Rosenschein. Negotiation and goal relaxation. Technical report, Computer Science Department, Hebrew University, Jerusalem, Israel, 1990. In preparation.