Forming Coalitions in the Face of Uncertain Rewards

Steven Ketchpel

Stanford University
Computer Science Department
Stanford CA, 94305
ketchpel@cs.stanford.edu

Abstract

When agents are in an environment where they can interact with each other, groups of agents may agree to work together for the benefit of all the members of the group. Finding these coalitions of agents and determining how the joint reward should be divided among them is a difficult problem. This problem is aggravated when the agents have different estimates of the value that the coalition will obtain. A "two agent auction" mechanism is suggested to complement an existing coalition formation algorithm for solving this problem.

1. The Problem

Given a set of agents with different abilities and different information, there may be many opportunities for cooperation among the agents that will benefit all. Even more likely is the chance that a *coalition* can form, a subset of the agents working together, benefiting each agent in the group perhaps at the expense of the community as a whole. An agent following the economic principle of rationality will attempt to form a coalition which will maximize its own utility. However, the other agents in these coalitions will have their own preferences, and a complicated cycle of dependencies emerges. Agents only want to commit to a coalition once all of the other agents have committed. The final division of the agents into coalitions should be *stable* in the sense that no subset of the agents could leave their current coalitions to form a new one yielding all of the agents in that new coalition a higher utility than they obtain from their previous coalitions.

For example, imagine there are a number of people interested in starting new hi-tech companies. There are many possible combinations of people that could work together, but getting them to commit to form a new company is difficult. A scientist with a hot new product idea doesn't want to commit unless a company has the necessary start-up capital. But financial backers typically require a thorough evaluation of the product and prefer a company president who has clout in the industry. The president may have reservations about working with certain financial officers, and so on. Even after the involved parties do agree to work together, bargaining over how to share the profits can reveal diverging perceptions about the relative importance of the different contributors. The coalition formation process, in this context, would describe which people should work together to start new companies and would also suggest a way to divide the profits among the partners. Stability in this scenario would mean that it would be unprofitable for one company to hire away workers from another, and there is no incentive for workers to get together (possibly with people from other companies) to start a new company.

To evaluate a system formally, the agents $a_1,\ldots,a_N$ are divided into a partition P containing coalitions $C_1,\ldots,C_M$ such that every agent is a member of exactly one coalition. The payoff to an agent is a function u(P, a) of both the partition and the agent. For P to be stable, there must not be any other partition P' forming coalitions $C'_1,\ldots, C'_M$ such that $\exists C'_i \in P' \; \forall a_j \in C'_i \;\; u(P', a_j) > u(P, a_j)$. If there were such a C'_i, the agents of that coalition would desert their current coalitions and form C'_i.

Determining how to divide the utility among the agents in the coalition is a problem that has received some attention in both game theory and distributed AI. A summary of the related research appears in Section 2. Many of these sources make the assumption that the value of any coalition is common knowledge. In game theoretic terms, there is a valuation function $v: 2^A \rightarrow \Re$, which takes any possible subset of the agent pool A, and returns a real value representing the utility which is split among the members of the coalition. For the sake of simplicity, we assume that this utility is paid by an entity outside the system of agents, and that none of the agents have any inherent interest in achieving the goals, beyond merely fulfilling the contract to receive payment. The main

The research was partially supported by a National Defense Science and Engineering Graduate Fellowship. The ideas contained within do not necessarily reflect the position or the policy of the Government and no official endorsement should be inferred.

contribution of this paper is to examine the case where the agents do not have access to this function, but instead have different expectations about the value.

Section 3 analyzes one of the most widely used division mechanisms, the Shapley value, and its inherent problems. Section 4 proposes an alternative approach which does not make the common knowledge assumption. The "Two Agent Auction" mechanism and the properties proved about it constitute the original research contribution. Section 5 concludes with directions for further research.

2. Related Research

At the 1993 European Workshop on "Modeling Autonomous Agents in a Multi-Agent World", three papers on coalition formation were presented [(Ketchpel 1993), (Shechory & Kraus 1993), (Zlotkin & Rosenschein 1993)]. The last two assumed super-additive domains in which adding an additional agent to a coalition can never reduce the utility of that coalition. Zlotkin and Rosenschein (Zlotkin & Rosenschein 1993) made the further stipulation that utility was not directly transferable between agents. All three papers assumed that the agents had common knowledge of the value function of the game and advocated the use of the Shapley value to divide the utility among the members of the coalition.

There is another body of literature in economics which addresses the division of goods or costs among the members of a society. Raiffa includes a chapter in his book (Raiffa 1982) on fair division and includes an analysis where the involved parties place different values on the goods to be divided. Ephrati and Rosenschein (Ephrati & Rosenschein 1991) use another device from economics known as the Clarke tax to allocate costs among multiple agents deciding among alternatives, charging each agent only in proportion to the amount it changed the group decision. The WALRAS system (Wellman 1993) uses a market scheme to reach an equilibrium among buyers and sellers of a commodity in the context of distributed action. However, none of these works analyzes the possibility of collusion by a coalition. This paper attempts to unite these two strands of research.

3. The Shapley Value and its Problems

The function u(P, a) determines the amount of utility that agent a receives from its membership in its coalition in P. It is assumed that the distribution is efficient and no utility is lost in the division, so $\sum_{a \in C} u(P, a) = \nu(C)$. There have been a number of suggestions for such a distribution function u(P, a). One of the earliest and most widely used is due to Shapley (Shapley 1953). The Shapley value is calculated by looking at each of the different dynamics that could lead to the coalition under consideration. Agents either "found" a coalition if they are the initial member, or else join a coalition founded by another member. The permutations of the members in the coalition is the set of formation dynamics. Each permutation describes an order in which the coalition could have been formed. Each agent adds value to a given formation process based on the marginal utility contributed by that agent. For example, if agent A is joining agents B, C, and D, and ν(ABCD) = 100 and ν(BCD) = 60, then A's marginal contribution under this formation ordering is ν(ABCD) – ν(BCD) = 40. If agent A joins a coalition started by B, and they are subsequently joined by agents C and D, A's marginal contribution is ν(AB) – ν(B). There are 22 other permutations that also might lead to the final coalition ABCD. By averaging A's marginal contribution across all the different formation possibilities, A's Shapley value is obtained. The underlying assumption is that all of the different formation processes are equally likely and, therefore, the marginal contributions for each formation are weighted equally. This calculation ensures that the sum of the Shapley values for all of the members of the coalition will be exactly the coalition's combined utility.

The Shapley value has several disadvantages. First, the most efficient known calculation is exponential, though efficient means to calculate the expectation of the Shapley value over a large number of interactions are known. (Zlotkin & Rosenschein 1994). Second, it assumes common knowledge of the value that the coalition will obtain if it works as a unit. In more realistic assessments, each agent might have a different expectation for the value of the collaboration.

To address these uncertainties more realistically, the value function should be dependent on which agent is performing the determination. That is, for two agents A and B, ν_A(AB), A's estimate of the value of coalition AB is not necessarily equal to B's estimate ν_B(AB), and both of these values may differ from the utility that will actually result from the coalition, which is denoted ν(AB) (and is the same ν(AB) used above). The potential disparity between these values (the actual utility and the various agents' estimates of it) opens up a further problem. One agent may overestimate the value, and promise its potential coalition partner a "share" of utility larger than the total obtained by the whole coalition. When the obtained utility fails to meet the rosy predictions of the optimistic agent, who is penalized?

4. Coalition Formation Using a Two Agent Auction

The problem that we are attempting to solve is two-fold: first, to determine coalitions of agents that will work together; second, to decide how to reward the agents, that is, what payment each agent will receive. These problems are complicated because the search space is very large (an exponential number of coalitions) and there are many dependencies among the decisions. For example, an agent's offer to join a coalition may depend on the agents already in the coalition, the amount of the offer, offers from other coalitions, and the future prospects of this coalition's merging with other coalitions. Finally, the agents may have different perceptions about the value of collaboration and their respective contributions to the group's outcome. The solution that we outline simplifies the problem along several dimensions, which we hope to address in future work.

The basic model that we assume is an economic one of rational agents entering into contracts that specify guaranteed payments. The agents may have different bargaining power due to their relative contributions to coalitions, but we assume that they all play symmetric roles in the bargaining process. The prescribed process consists of the following steps:

1. Agents exchange initial offers to other available agents. These offers will lead to a possible agreement and contract among the agents.
2. Agents evaluate the offers they received, and rank them in order of preference, based on their expected profit.
3. Using these preference orderings, the agents attempt to pair off into coalitions of size 2 with the most attractive potential partners.
4. The newly formed pairs enter a "two agent auction" that makes one agent the manager, bearing the risk and given the opportunity to bargain on behalf of the pair in future negotiations. The non-managing agent receives a fixed payment for its role in the coalition. The final agreement price is a function of the initial offers and the agents' valuations of the collaborative effort.
5. The process repeats, with the pairs formed in one round playing the role of individual agents in the next.

4.1. A Coalition Formation Algorithm

In previous work (Ketchpel 1993), we noted that the coalition formation problem is related to the stable marriage problem (Gusfield & Irving 1989). In the stable marriage problem, an equal number of men and women seek mates. Each participant has a preference ordering among the candidates, and a stable matching is generated when each man is paired with a woman and there is no blocking pair of a man and woman that prefer to be paired with each other to being paired with their current partners. A stable matching may be found for any instance of the problem in time $O(n^2)$ where n is the number of people involved.

The coalition formation process for coalitions of size 2 is equivalent to a variant of the stable marriage problem known as the stable roommate problem with unacceptable partners. In the stable roommate problem, the two classes of men and women are conflated to a single class, agents. When unacceptable partners are allowed, an agent prefers being unpaired to being paired with certain other agents. A pairing which matches any agent with an unacceptable partner is inherently unstable. Centralized versions of the stable roommate problem with unacceptable partners find stable matchings (when they exist) in time $O(n^2)$.

However, in a setting of autonomous, distrustful agents, a centralized algorithm is not a viable solution. In (Ketchpel 1993), a decentralized alternative is proposed. The modified algorithm is a greedy process where each agent proceeds down its preference list extending an offer to the top agent it hasn't previously asked, accepting offers that improve its utility, and rejecting all others. At the end of a round, all of the pairs form proto-coalitions, which may join other proto-coalitions in the future. They select one of the members to act as the head of the coalition. In the subsequent rounds, the process repeats, with each coalition head extending offers to the heads of other coalitions and to agents that have not yet been paired. The process repeats until no new associations are formed. The algorithm takes time $O(n^3)$ for n agents. Although stability is not guaranteed, an agent will never settle for a less desirable coalition partner unless all of the better alternatives (taking the previous rounds of formation as given) have turned it down once already. Even if the other possible partners have turned it down in the past, they may later be willing to accept such a coalition. The agent will never approach these possible partners again, so unstable pairings may form. For a more complete description and complexity analysis, see (Ketchpel 1993).

4.2. The Two Agent Auction

One mechanism to solve the division of utility in the face of uncertainty is to assign one of the agents responsibility for managing the group actions. The manager is required to meet the offers that it extended to the various coalition members, even if the coalition's actual utility were less than expected. In exchange for undertaking this risk, the managing agent would receive all of the utility accruing to the coalition, and would earn a profit if this amount were greater than the salaries it paid. Also, as the manager, it has the authority to negotiate on behalf of the group to form larger coalitions.

The algorithm described in Section 4.1 has the property that each of the proto-coalitions has exactly two entities (which may be agents or coalitions). Therefore, each of the auctions occurs between two agents, the managers of the coalitions that are merging. The two managing agents A and B begin the bargaining process using the initial offers that they extended to each other when the preference lists for the previous step were made. These offers will not necessarily add up to either agent's estimate of $v(AB)$, nor need they total the actual $v(AB)$ value. The offers are adjusted according to the method described below and summarized in Figure 1. The two agents are guaranteed to converge on an agreeable value. The non-managing agent gets this agreed value, regardless of the actual utility of the coalition. The managing agent receives the balance of the utility obtained by the group. We use O(A, B) to represent the amount of the initial offer which agent A extended to agent B; similarly, O(B, A) is B's initial offer to A.

In selecting the agent to be the manager, there are four cases that may occur:

1. Both agents A and B want to be the manager, based on the offers and their beliefs about the actual value of the collaboration. So, $v_A(AB) - O(A, B) > O(B, A)$ and

$v_B(AB) - O(B, A) > O(A, B)$. The agents reach agreement through an ascending auction.

2. Agent A wants to be the manager, and agent B is happy to agree. So, $v_A(AB) - O(A, B) > O(B, A)$ and $v_B(AB) - O(B, A) \leq O(A, B)$. In this case, A is selected to be the manager.

3. Symmetric to 2, with B wanting to be the manager.

4. Neither agent wants to be the manager, because both expect better payoffs if the other agent is the manager. So, $v_A(AB) - O(A, B) \leq O(B, A)$ and $v_B(AB) - O(B, A) \leq O(A, B)$. The agents reach agreement by entering a descending auction.

In the first case, there needs to be further negotiation over who will manage the contract. To settle the difference, both agents incrementally increase their offers to the other coalition agent until one or the other is willing to forgo the opportunity to be the manager. In essence, the two agents are "bidding" for the right to manage the contract.

```
BEGIN.
  k := 0.  /*k is number of rounds of negotiation conducted*/
  δ := 1.  /*δ is "precision" of negotiation*/
  IF  vA(AB) - O(A, B) > O(B, A)
        AND vB(AB) - O(B, A) > O(A, B)
     I := +1.    /*Reduce Case 1 to 2 or 3*/
     WHILE ((vA(AB) -  (O(A, B) + I*k*δ) > (O(B, A) + I*k*δ)
           AND vB(AB) -  (O(B, A) +I*k*δ)> (O(A, B)+ I*k*δ))
              k := k + 1.
     END-WHILE.
  ELSE
     IF vA(AB) - O(A, B) ≤ O(B, A)
           AND vB(AB) - O(B, A) ≤ O(A, B))
        I := -1.    /*Reduce Case 4 to 2 or 3*/
        WHILE ((vA(AB) -  (O(A,B) + I*k*δ)<(O(B,A)+ I*k*δ)
           AND vB(AB) -  (O(B,A) + I*k*δ) ≤ O(A,B) + I*k*δ))
              k := k +1.
        END-WHILE.
     END-IF.
  END-IF.

  IF vA(AB) -  (O(A, B) + I*k*δ) ≥ O(B, A) + I*k*δ
     A is manager, B gets O(A, B).+ I*k*δ      /*Case 2*/
  ELSE
     B is manager, A gets O(B, A)+  I*k*δ      /*Case 3*/
  END-IF.
END.
```

Figure 1: Algorithm for selecting manager & determining utility division

In the ascending auction called for in the first case, at each iteration of the WHILE loop in Figure 1, both agents increase their offers by δ. The bidding stops when either agent finds that the "opposing" agent (although they are coalition partners, they are competing with each other to maximize individual shares of the joint gain) has extended an offer that is greater than it would expect if it managed the contract. Note that there is some asymmetry in the roles of the agents. In one case the test is a strict inequality, while in the other case, the test is less than or equal to. We arbitrarily select the agent that initiates the proposal to be agent A.

In the fourth case in which neither agent wants to be the manager, the agents enter an auction situation similar to case 1, but instead of incrementing their offers, they decrement them. At some point one of the agents will decide that with this new lower offer, it is better to accept the managing role than the small amount just promised by the other agent. This agent is made the manager, and its last offer is considered the agreement value.

As an example, assume that agents A and B have agreed to form a coalition, and are trying to determine the distribution of the utility from the joint effort. Agent A expects that the value of the outcome will be 100, so $v_A(AB) = 100$. Agent A realizes that agent B is doing a larger share of the work, so is willing to offer agent B a larger share of the utility, in this case, $O(A, B) = 60$. Agent B is more pessimistic about the expected outcome of their joint effort, expecting only 80 units of utility $v_B(AB) = 80$. Agent B thinks that agent A's contribution is minimal and is only willing to give agent A 15 units, $O(B, A) = 15$. The case analysis outlined above shows that this example falls in the first case, and both agents A and B want to manage the contract. Agent A's expected profit if it is the manager is 40 ($v_A(AB) - O(A, B)$); if A accepts B's offer, A will only obtain 15. Agent B carries out a similar analysis and sees that its expected return of 65 if it manages the contract ($v_B(AB) - O(B, A)$) exceeds A's offer of 60. At this point, the negotiation enters the stage of incrementally increasing offers. The progress of these iterative offers is shown in Figure 2. At round 3, B determines that it expects to get more if it allows A to manage the contract, so A is obligated to pay B 63 units of utility when B accomplishes its share of the work, and agent A will get the actual amount $v(AB)$. If this amount is less than 63, A still must pay B the promised 63 units. If $v(AB)$ is less than 81, then A would have been better off accepting B's offer of 18, rather than receiving $v(AB)$ while paying agent B 63.

	A's Expected Value if:		B's Expected Value if:	
	A manages	B manages	B manages	A manages
k	$v_A(AB) - O(A, B) - k * \delta$	$O(B, A) + k * \delta$	$v_B(AB) - O(B, A) - k * \delta$	$O(A, B) + k * \delta$
1	39	16	64	61
2	38	17	63	62
3	37	18	62	63

Figure 2: Sequence of offers between agents

4.3 Analysis of the Two Agent Auction

Although the negotiation is described above in an incremental process, the result is deterministic. The agent with the higher estimate of $v(AB)$ always becomes the manager, as is shown in Figure 3. Moreover, Theorem 2 in Figure 4 shows that the agreement price is also determined by the initial offers and valuations. If the agents are willing to share their estimates of $v(AB)$ with their initial offers, they can directly calculate the differences between the evaluations of the agents' contributions and determine which agent should be the manager and what the final offer to the non-managing agent should be. If the $v(AB)$ estimates are not shared, the iterative method described above will yield the same result, though the manager's estimate of $v(AB)$ will never become public knowledge. The choice of incremental versus direct calculation is dependent on the domain, and the tradeoff between the benefit of privacy of information against the cost of more communication.

Theorem 1: Between two agents A and B, the one with the higher valuation of $v(AB)$ will always win the managing role.

The auction stops after k rounds, when either:

1) $O(B, A) + k*\delta > v_A(AB) - (O(A, B) + k*\delta)$; B manages

or

2) $O(A, B) + k*\delta \geq v_B(AB) - (O(B, A) + k*\delta)$; A manages

If (1) is the reason for stopping,

(1a) $O(B, A) + k*\delta > v_A(AB) - (O(A, B) + k*\delta)$

and (1b) $O(A, B) + k*\delta < v_B(AB) - (O(B, A) + k*\delta)$

Adding $k*\delta$ to both sides of 1a and 1b,

(1a') $O(B, A) + 2*k*\delta > v_A(AB) - O(A, B)$

(1b') $O(A, B) + 2*k*\delta < v_B(AB) - O(B, A)$

Adding O(A, B) to both sides of (1a') and O(B, A) to both sides of (1b')

(1a") $O(A, B) + O(B, A) + 2*k*\delta > v_A(AB)$

(1b") $O(A, B) + O(B, A) + 2*k*\delta < v_B(AB)$

By transitivity of 1a" and 1b"

$v_A(AB) < v_B(AB)$, and in (1), B is the manager

If (2) is the reason for stopping,

(2a) $O(B, A) + k*\delta \leq v_A(AB) - (O(A, B) + k*\delta)$

and (2b) $O(A, B) + k*\delta > v_B(AB) - (O(B, A) + k*\delta)$

Proof proceeds as above, replacing the strict inequality with non-strict inequality, yielding,

$v_B(AB) \leq v_A(AB)$, and in case 2, A is the manager

So, in both cases, the agent with the higher estimate of $v(AB)$ is the manager.

Figure 3: Agent with higher estimate of $v(AB)$ is manager

Theorem 2: The agreement price (AP) will be within δ of $\frac{O(M,N) + v_N(MN) - O(N,M)}{2}$, where M is the manager, N is the other (non-managing) agent.

The auction will stop in round k when

$$O(M, N) + k*d \geq v_N(MN) - (O(N,M) + k*\delta)$$

$$O(M, N) + 2*k*\delta \geq v_N(MN) - O(N,M)$$

$$2*k*\delta \geq v_N(MN) - O(N,M) - O(M, N)$$

$$k = \left\lceil \frac{v_N(MN) - O(N,M) - O(M,N)}{2*\delta} \right\rceil$$

The offer after k rounds of negotiation is $O(M, N) + k*\delta$.

$$AP = O(M,N) + \left\lceil \frac{v_N(MN) - O(N,M) - O(M,N)}{2*\delta} \right\rceil * \delta$$

$$O(M,N) + \frac{v_N(MN) - O(N,M) - O(M,N)}{2*\delta} * \delta \leq AP, \text{ and}$$

$$AP < O(M,N) + \left(\frac{v_N(MN) - O(N,M) - O(M,N)}{2*\delta} + 1 \right) * \delta$$

$$\frac{O(M,N) + v_N(MN) - O(N,M)}{2} \leq AP, \text{ and}$$

$$AP < \frac{O(M,N) + v_N(MN) - O(N,M)}{2} + \delta.$$

So, AP is within δ of $\frac{O(M,N) + v_N(MN) - O(N,M)}{2}$

Figure 4: Agreement price is function of offers and $v(AB)$ estimates

4.4. Selecting Initial Offers

The agreement price that is reached is a function of the initial offers and the estimates of $v(AB)$ as Figure 4 shows. From the final value, it appears that both agents will extend initial offers of 0. The agreement price increases with O(M, N), the initial offer of the manager to the non-manager. Therefore, an initial offer of 0 would minimize the agreement price with respect to this variable. Likewise, the agreement price decreases as the offer of the non-manager to the manager increases, so an initial offer of 0 would maximize the agreement price. However, this analysis is too simplified. The initial offers play a second role in the coalition formation process, namely determining the preference lists. Therefore, the agents need to extend sufficiently high offers to each other to ensure that the other agent will agree to form a coalition. The auction mechanism (and the desire to minimize the initial offers) is only needed after two agents have agreed to form a coalition. In order to select an optimal initial offer, agents need to take into account expectations about the offers that other agents will make. This game theoretic analysis requires more complicated machinery that is beyond the scope of this paper, and is grounds for future research.

5. Conclusion and Future Research

This paper has described a utility distribution mechanism designed to perform in situations where there is uncertainty in the utility that a coalition obtains. The two agent auction is shown to have certain properties: first, the agent valuing the collaboration more highly is always selected manager; second, the agreement price is a deterministic function of the agent's initial offers and estimates of the value of collaboration.

A logical next step would be to further develop the game theoretic aspects of this model. First, the agents need additional information to strategically make their initial offers. Second, in the final stages of the two agent auction, an agent might want to report a higher estimate of v(AB) than it truly holds, knowing that if the other agent has a higher value still, this deception will raise the agreement price.

Agents may also want to form coalitions that temporarily decrease their expected utility in the short run, with the expectation that this coalition will be able to improve its position by bargaining in future rounds. A related analysis (Aumann & Myerson 1988) looks only at the case of super-additive environments. A further interesting problem is to see if agents can infer useful information based on the preferences revealed by the offers, then act strategically with this new knowledge.

Acknowledgments

This work would not have been possible without the guidance and support of the numerous people who made considerable contributions. Steven Brenner, Barbara Grosz, Joe Halpern, Victor Milenkovic, Jeff Rosenschein, Dan Roth, Stuart Shieber, Yoav Shoham, the Logic Group at Stanford, and anonymous reviewers deserve mention for thought-provoking discussions and helpful suggestions.

References

Aumann, Robert J., and Myerson, Roger B. 1988. Endogenous formation of links between players and of coalitions: an application of the Shapley value. In Roth, Alvin E., ed., *The Shapley Value*, chapter 12, pp. 175–191. Cambridge University Press.

Ephrati, Eithan and Rosenschein, Jeffrey S. 1991. The Clarke Tax as a consensus mechanism among automated agents. In Proceedings of the Ninth National Conference on Artificial Intelligence, 173–178. Menlo Park, CA: AAAI Press.

Gusfield, Dan and Irving, Robert W. 1989. *The Stable Marriage Problem: Structure and Algorithms*. Cambridge, MA: MIT Press.

Ketchpel, Steven P. 1993. Coalition Formation Among Autonomous Agents. In Pre-Proceedings of the 5th European Workshop on "Modeling Autonomous Agents in a Multi-Agent World" Supplement. Neuchâtel, Switzerland: Institut d'Informatique et Intelligence artificielle, Université de Neuchâtel.

Raiffa, Howard. 1982. *The Art and Science of Negotiation*. Cambridge, MA: Belknap Press.

Shapley, Lloyd. 1953. A Value for N-Person Games. In Kuhn, H. W. and Tucker. A. W. eds. *Contributions to the Theory of Games*. Princeton University Press.

Shechory, On and Kraus, Sarit. 1993. Coalition Formation Among Autonomous Agents: Strategies and Complexity. In Pre-Proceedings of the 5th European Workshop on "Modeling Autonomous Agents in a Multi-Agent World". Neuchâtel, Switzerland: Institut d'Informatique et Intelligence artificielle, Université de Neuchâtel.

Wellman, Michael P. 1993. A Market-Oriented Programming Environment and its Application to Distributed Multicommodity Flow Problems. *Journal of Artificial Intelligence Research.* 1: 51–92.

Zlotkin, Gilad and Rosenschein, Jeffrey S. 1994. Coalition, Cryptography, and Stability: Mechanisms for Coalition Formation in Task Oriented Domains. In Proceedings of the Twelfth National Conference on Artificial Intelligence. To appear.

Zlotkin, Gilad and Rosenschein, Jeffrey S. 1993. One, Two, Many: Coalitions in Multi-agent Systems. In Pre-Proceedings of the 5th European Workshop on "Modeling Autonomous Agents in a Multi-Agent World". Neuchâtel, Switzerland: Institut d'Informatique et Intelligence artificielle, Université de Neuchâtel.