Passive Threats among Agents in State Oriented Domains

Yair B. Weinberger and **Jeffrey S. Rosenschein**[1]

Abstract.

Previous work in multiagent systems has used tools from game theory to analyze negotiation among automated agents in cooperative domains. Rosenschein and Zlotkin, using these tools, provided a general mechanism for two-agent negotiation in an isolated State Oriented Domain (SOD) encounter, and also provided a classification that divided these encounters into four basic types. Other multiagent systems work considered the notion of threats during negotiation, but in order to do so introduced additional assumptions on the domain.

This paper presents a new model of threats among negotiating agents in State Oriented Domains that requires no additional domain-specific assumptions. We assume that agents may use a "threat of passivity" against other agents — in other words, threatening to remain inactive (and not exploit existing cooperative opportunities), forcing both agents to satisfy their goals on their own (serially). The possibility of this negotiation threat adds interesting complexity to the four basic SOD encounter types, further subdividing them into additional types of encounter. We analyze these new encounter types that arise when there is the possibility of "passive threats", providing a thorough characterization of their properties.

1 INTRODUCTION

Negotiation plays an important role in encounters between self-interested agents, and has been widely studied in multiagent systems (MAS) research [9, 3]. Game theory tools have been particularly helpful in understanding the formal properties of negotiation, in various kinds of encounters. The negotiation literature specifies properties that might be exhibited by negotiation protocols and by the deals that those protocols produce [8]. For example, the property of "fairness" of a deal [12]—that it maximizes the product of agents' utilities—is sometimes enforced by a given negotiation mechanism.

Assuming that agents have agreed on certain ground rules (e.g., what constitutes the negotiation set), there is still the possibility of manipulation of the negotiation through various techniques. Prior research [9] has discussed the question of how an agent—given the same basic encounter rules—might improve his own outcome at the expense of the other agent. For example [11, 6], one can show that making up goals or declaring goals to have a false worth (baseline utility) can achieve a better deal. Techniques that cause a false computation of u also cause a false computation of the product of utilities.

1.1 Utility Improvement through Threats

In this paper, we take a different approach to the question of how one agent can improve a negotiation outcome for itself. Instead of lying about its goals, or the worth of its goals, a *threat* causes an actual change in the computation of the other agent's utility. In order for the threatened agent to maximize its utility, it will agree to a deal that gives the threatening agent higher utility than the utility the threatening agent could have achieved in a regular negotiation.

Threats are based on evaluating what the other agent will do if the negotiation fails, and providing him with a slightly better option. Being rational, the other agent will accept the threat and agree to the deal, even though the deal does not satisfy the product maximizing condition mentioned above.[2]

The notion of threats in multiagent negotiation has appeared in a variety of prior work [2, 10, 3, 7, 4, 5]. Much of this research deals with the formal definition of threats, and how to add threats to the formal protocols of negotiation. Other papers present several kinds of threats, and deal with issues like the reliability of threats and the ability of agents to carry out the threats. The threats considered in this prior work posit additional assumptions on the domain, including limitations on the computational power of agents, assumptions that the agents exist in a society that remembers the reputation of each one, or even assumptions that agents live in a physical world, and that they can physically harm one another.

The "threats of passivity" that we explore here can in principle be applied to any State Oriented Domain that fulfills the assumptions in [9]. Thus, they are applicable to real-world domains where negotiation among intelligent agents establishes allocations of resources or of effort—such as Web services, general e-commerce, or information management.

1.2 Definitions

In this paper, we view threats in a domain-independent fashion, focusing on the State Oriented Domains of [9]. We add another course of action for the agents to pursue—in other words, another row to the matrix of a game in normal form—which we refer to as a "passive threat". Adding this course of action does not require any underlying assumption whatsoever beyond those that are used in the model presented in [9].

The passive threat has the following form: "If you do not agree to the deal that I am offering you, I will stand aside and let you satisfy your goal alone, and after you finish I will satisfy my goal, while keeping your goal satisfied. If you also choose to stand aside, I will satisfy my goal in a way that will be much worse for you." This course of action, while requiring no specific assumptions on the State Oriented Domain, adds some interesting complexity to the four basic encounter types presented by Rosenschein and Zlotkin, and engenders different results in some encounters.

We adopt the basic notation of [12]. Limitations of space prevent us from providing a full description of notation; however, the following terms are important for our own analysis:

[1] School of Engineering and Computer Science, The Hebrew University of Jerusalem, Jerusalem, Israel; email: {wein, jeff}@cs.huji.ac.il

[2] Experiments with humans [1] arrive at different conclusions. People will often forgo a deal if they feel it treats them unfairly—even if declining the deal leaves them even worse off. We are, however, interested in the formal properties of rational agents, not in a descriptive analysis.

Definition 1 *A State Oriented Domain (SOD) is a tuple* $< S, A, \Im, c >$ *where:*

1. S *is the set of all possible world states;*
2. $A = \{A_1, A_2, ..., A_n\}$ *is an ordered list of agents;*
3. $\Im$ *is the set of all the possible joint (i.e., n-agent) plans. A joint plan* $J \in \Im$ *moves the world from one state in* S *to another. The actions taken by agent* k *are called* k*'s role in* J*, and will be written as* J_k*. We can also write* J *as* $(J_1, J_2, ..., J_n)$*;*
4. c *is a function* $c : J \rightarrow (\mathbb{R}^+)^n$*. For each joint plan* $J \in \Im$*,* $c(J)$ *is a vector of* n *positive real numbers, the cost of each agent's role in the joint plan.* $c(J)_i$ *is the* i*-th element of the cost vector, i.e., it is the cost of the* i*-th role in* J*. If an agent plays no role in* J*, his cost is 0.*

Definition 2 *An encounter within an SOD* $< S, A, \Im, c >$ *is a tuple* $< s, (G_1, G_2, ..., G_n) >$ *such that* $s \in S$ *is the initial state of the world, and for all* $k \in \{1, .., n\}$*,* G_k *is the set of all acceptable final states from* S *for agent* A_k*.* G_k *will also be called* A_k*'s goal.*

Definition 3 *Worth: Given an encounter in a two-agent SOD* $< s, (G_1, G_2) >$*, let* w_i *be the maximum expected cost that agent* i *is willing to pay in order to achieve his goal* G_i*.* w_i *will be called the worth of goal* G_i *to agent* i*. We will denote this enhanced encounter by* $< s, (G_1, G_2), (w_1, w_2) >$*.*

Definition 4 *Given an encounter* $< s, (G_1, G_2), (w_1, w_2) >$*, if* δ *is a deal, i.e., a mixed joint plan satisfying both agents' goals, then* $Utility_i(\delta)$ *is defined to be* $w_i - cost_i(\delta)$*.*

Definition 5 *Let* J *be a joint plan of two agents,* A_1 *and* A_2*. Then a "fair" deal satisfies* $J = \arg\max_{J \in NS}\{Utility_1(J) \cdot Utility_2(J)\}$*, i.e., it is the deal in the negotiation set (pareto optimal, individual rational agreements) that maximizes the product of the two agents' utilities.*

The following interaction types were presented in [9].
1) In a symmetric cooperative situation there exists a deal in the negotiation set that is preferred by both agents over achieving their goals alone. Here, both agents welcome the existence of the other agent.
2) In a symmetric compromise situation there are individual rational deals for both agents. However, both agents would prefer to be alone in the world, and to accomplish their goals alone. Since each agent is forced to cope with the presence of the other, he would prefer to agree on a reasonable deal. All of the deals in the negotiation set are better for both agents than leaving the world in its initial state s.
3) In a non-symmetric cooperative/compromise situation, one agent views the interaction as cooperative (he welcomes the existence of the other agent), while the second agent views the interaction as compromise (he would prefer to be alone in the world).
4) In a conflict situation the negotiation set is empty—no individual rational deals exist.

2 EXAMINING ENCOUNTER TYPES

2.1 Symmetric Cooperative Encounters

In a cooperative encounter, when an agent carries out his task, he also helps the other agent. Figure 1 gives an example of this kind of encounter. Alone in the world, agent A_1 can achieve his goal with cost 4, and A_2 can achieve his goal with cost 2. There is a joint plan with a cost of 2 to each agent. But what happens if A_2 refuses to accept the joint plan? The negotiation fails, and A_1 will carry out his task alone. As a side effect of carrying out his task, A_1 also carries out A_2's task. A_2's goal is now achieved with a cost of 0.

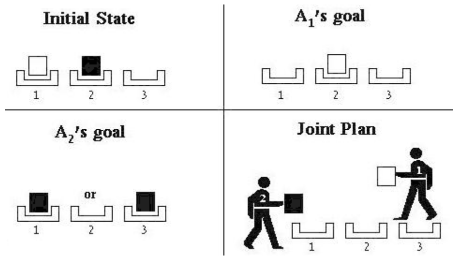

Figure 1. A Symmetric Cooperative Encounter

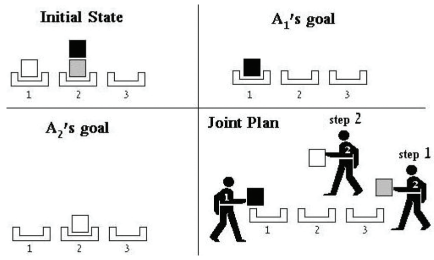

Figure 2. A Cooperative Encounter that is Beneficial to Both Agents

Claim 1 *Given a cooperative encounter in a two-agent SOD* $\langle s, (G_1, G_2) \rangle$*, for every* $g \in G_1$*,* $c(s \rightarrow^2 G_2) \geq c(g \rightarrow^2 G_2)$

Proof. Follows from the definitions. ■

Note that this claim does not imply that an agent should never agree to a joint plan in a cooperative encounter. The example in Figure 2 shows why. As we can see, the best joint plan that satisfies the product maximizing condition on utility is the one in the picture, where each agent takes each role with a probability of 0.5. This plan gives an expected cost of 3 to each agent. Figure 3 shows the world after A_2's best plan if A_1 does not agree to the joint plan, but instead waits for A_2 to carry out his task. A_1 still must use 4 pickup/putdown operations in order to carry out his goal. Of course, it would have been better for him to agree to the joint plan.

2.2 Multiple Symmetric Cooperative Encounters

We can, in fact, distinguish between several types of symmetric cooperative encounters; we will now formally define those types. Below, if one agent is denoted as A_i, then the other will be denoted as A_j.

Definition 6 *Let* $\langle s, (G_1, G_2) \rangle$ *be a symmetric cooperative encounter in a two-agent SOD.* $\langle s, (G_1, G_2) \rangle$ *is said to be a* full cooperative encounter *if there exists a joint plan J such that:*

$$\forall i \in \{1, 2\}, \forall g \in G_i \quad \text{s.t. } g \text{ is reachable by agent } i \text{ alone,}$$
$$\text{it is true that} \quad c_j(J) \leq c(g \rightarrow^j G_j).$$

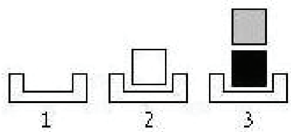

Figure 3. The World After A_2 Works Alone

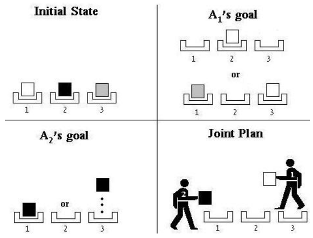

Figure 4. An Example of a Semi-Cooperative Encounter

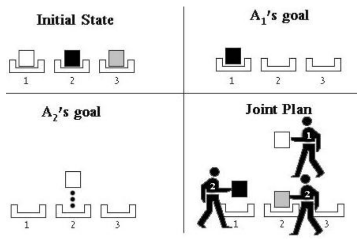

Figure 5. An Example of a Cooperative/Compromise Encounter

The encounter in Figure 2 is an example of a full cooperative encounter. It is beneficial for both agents to take part in the joint plan. It is not good for any agent to wait for the other to work alone, since taking part in the joint plan results in better utility for each.

Definition 7 *Let $\langle s, (G_1, G_2)\rangle$ be a symmetric cooperative encounter in a two-agent SOD that is not a full cooperative encounter. $\langle s, (G_1, G_2)\rangle$ is said to be a* semi-cooperative encounter *if there exists a joint plan J such that:*

$$\forall i \in \{1, 2\}, \exists g \in G_i \text{ s.t. } g \text{ is reachable by agent } i \text{ alone and } c_j(J) \leq c(g \rightarrow^j G_j).$$

Figure 4 shows an example of a semi-cooperative encounter. A_1 would like to use the joint plan in order to achieve his first goal with a cost of 2. Alone, he can achieve the first goal with a cost of 4, and the second goal with a cost of 6. If he will pursue his first goal, he will achieve A_2's goal along the way. As we will soon see, A_1 can leverage his second goal in order to make A_2 use the joint plan.

The definition of full cooperative encounters implies that each agent can "force" (using a threat that we will discuss below) the other agent to take part in the joint plan. We now formalize the above statement and prove it.

Claim 2 *Let $< s, (G_1, G_2) >$ be a full cooperative encounter with a joint plan J. Then every agent can act such that the best strategy for the other agent will be to take part in the joint plan.*

Proof. Without loss of generality, assume A_1 tries to force A_2 to take part in the joint plan. Let g be the encounter from the definition of full cooperative encounters. A_1 can *threaten* that if A_2 will not take part in the joint plan, then A_1 will reach alone the state g (which satisfies his goal). A_2 will then have to pay more to achieve his goal than he would pay with the joint plan. This is exactly the kind of threat that we presented in Section 1.2; it requires no further assumptions on the domain, following directly from the definition. ■

Note that we use the word "strategy" in order to represent the chosen action from a set of possible actions. As we said before, the possible actions are those imported from [9], and the added action of waiting. We now define the last type of "cooperative" encounter, which, in a sense, is not really cooperative.

Definition 8 *Let $< s, (G_1, G_2) >$ be a symmetric cooperative encounter in a two-agent SOD. $< s, (G_1, G_2) >$ is said to be a* dummy cooperative encounter *if:*

$$\exists i \text{ s.t. } \forall J \in NS, \forall g \in G_j \text{ it is true that } c_i(J) \geq c(g \rightarrow^i G_i).$$

The meaning of this definition is that one of the agents always prefers to let the other agent work alone, and only then finish what is left to achieve his own goal. The example that we saw above in Section 2.1 was an example of a dummy cooperative encounter. A_2 had no reason to cooperate.

2.3 Non-symmetric Cooperative/Compromise

In cooperative/compromise encounters, one agent views the interaction as cooperative, while the other agent views it as compromise. The first agent welcomes the existence of the other, while the second would prefer to be alone. We now differentiate among several types of cooperative/compromise encounters. Let us begin with an example of a cooperative/compromise encounter, as shown in Figure 5.

In this example, if both agents are to achieve their goals, the best joint plan has two roles; one role (A_1's role in the picture) has a cost of 2 pickup/putdown operations, and the other role (A_2's) has a cost of 4 pickup/putdown operations. Assuming both agents have the same worth to their goals, the deal that satisfies the product maximization condition is the one where each agent takes the first role with a probability of 0.5, and the second role with a probability of 0.5. This deal results in expected work of 3 for each agent.

Compare this joint plan with the encounter where each agent is alone in the world. This is better for A_1, because if he were alone in the world, he would have to do 4 pickup/putdown operations in order to achieve his goal. However, it is worse for A_2, because if he were alone in the world he would only have to do 2 pickup/putdown operations, instead of the expected 3 he has to do in the joint plan.

Let us analyze this encounter using passive threats; we want to know if agent A_2 must compromise in this encounter. What happens if agent A_2 refuses the deal as it is now, and offers a new deal, using the same joint plan but with different probabilities? He offers, that with a probability of ε, he will take the second role (the one that costs 4), and with a probability of $1 - \varepsilon$ he will take the first role in the joint plan (the role that costs 2). This offer still gives A_1 expected work of less than 4. The expected work of agent A_1 from this deal is $4 - 2\varepsilon$; thus, it is rational for agent A_1 to agree to the deal, and the expected work of agent A_2 from the deal is $2 + 2\varepsilon$.

In the above discussion, we did not limit our choice of ε. A_2 can choose $\varepsilon > 0$ as small as he wishes, and thus derive expected work as close to 2 as he wishes, and achieve his goal without compromising. However, is this threat always useful? The answer is no. If after agent A_1 satisfies his goal alone (and pays the higher cost), A_2 needs, in order to achieve its goal, to pay more than he would pay in the joint plan, then agent A_1 can threaten to pay the higher cost and achieve his goal alone, if A_2 will not agree to the joint plan. It would be rational for A_2 to agree to the joint plan.

As we saw before in cooperative encounters, we now need to check if A_1 has a way of forcing A_2 into compromising. We will formally distinguish between two types of compromise encounters: encounters where one of the agents can force the other agent into compromising, and encounters where the compromising agent does not really *have* to compromise.

Definition 9 *Let P-NS be J s.t. $J = \arg\max_{J \in NS}\{u_1(J) \cdot u_2(J)\}$*

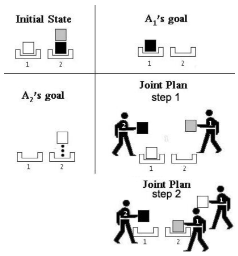

Figure 6. An Example of a Dummy Cooperative/Compromise Encounter

Claim 3 *Let $\langle s, (G_1, G_2)\rangle$ be a non-symmetric cooperative/compromise encounter in a two-agent SOD. Let A_1 be the cooperating agent, and let A_2 be the compromising agent. Then $\forall J \in$P-NS it is true that $c_2(J) \geq c(s \longrightarrow^2 G_2)$.*

Proof. Follows from the definitions. ■

2.3.1 Dummy Cooperative/Compromise Encounters

Definition 10 *Let $\langle s, (G_1, G_2)\rangle$ be a non-symmetric cooperative/compromise encounter in a two-agent SOD. Let A_1 be the cooperating agent, and let A_2 be the compromising agent. $\langle s, (G_1, G_2)\rangle$ is said to be a* dummy cooperative/compromise encounter *if: $\forall g \in G_1$ such that g is achievable alone by A_1, it is true that $\forall J \in$P-NS $c_2(J) \geq c(g \longrightarrow^2 G_1 \cap G_2)$, and $\forall J \in$P-NS $\exists J' \in NS$ such that $c_2(J) > c_2(J') \wedge c_1(J') \leq c(s \longrightarrow^1 G_1)$*

Figure 6 shows an example of a dummy cooperative/compromise encounter. The most simple example is when $\forall g \in G_1$, g is not achievable alone by A_1. In this case, the condition in the definition is trivially true (assuming that the worth is large enough).

Agent A_1 cannot achieve his goal alone. Agent A_2 can achieve his goal alone with a cost of 2. However, if both agents are to achieve their goals, the best joint plan gives each one expected work of 3. Assuming both worths are greater than 3, this is a non-symmetric cooperative/compromise encounter. Because agent A_1 cannot achieve his goal alone, this is a dummy cooperative/compromise encounter. Agent A_2 does not have to compromise, but can offer a deal where he takes the role that costs 2 (the role of agent 2 in the picture), and let A_1 take the role that costs 4. If A_1's worth is greater than 4, he will agree to the deal.

We will now divide the dummy cooperative/compromise encounter into three subtypes of encounters. We will see that one subtype of dummy cooperative/compromise encounter is really a *cooperative encounter*, in the sense that both agents welcome the existence of the other in the world. In the second subtype of these encounters, the compromising agent has to compromise, but still, the compromise is better for him than just using the deal that satisfies the product maximization condition. The third subtype is merely the thin line between the two other subtypes, where the compromising agent does not have to compromise, but he also does not welcome the existence of the other agent in the world.

Definition 11 *Let w_i denote the goal worth for A_i.*
Definition 12 *Let s be a state of the world. Let $c(s \longrightarrow^i G_i) \equiv w_i$ if there is no single agent plan $\mathbb{P}$ such that $\mathbb{P}$ starts in s and ends in $g \in G_1$.*

Definition 13 *Let $\langle s, (G_1, G_2)\rangle$ be a dummy cooperative/compromise encounter. Let A_1 be the cooperating agent, and let A_2 be the compromising agent. We will say that $\langle s, (G_1, G_2)\rangle$ is a* full dummy cooperative/compromise encounter *if: $\exists J \in NS$ such that $c_2(J) < c(s \longrightarrow^2 G_2)$ and $c_1(J) \leq c(s \longrightarrow^1 G_1)$.*

In a full dummy cooperative/compromise encounter, the compromising agent can benefit from the existence of the other in the world. He will not agree to the standard negotiation deal, but will offer the joint plan J from the definition. That joint plan will give him more than working alone—the other agent cannot force him into cooperation since it is a dummy cooperative/compromise encounter; and again, from the definition, it is rational for him to agree to this deal.

Definition 14 *Let $\langle s, (G_1, G_2)\rangle$ be a dummy cooperative/compromise encounter. Let A_1 be the cooperating agent, and let A_2 be the compromising agent. We will say that $\langle s, (G_1, G_2)\rangle$ is a* semi-dummy cooperative/compromise encounter *if: $\forall J \in NS$ such that $c_2(J) \leq c(s \longrightarrow^2 G_2)$ it is true that $c_1(J) \geq c(s \longrightarrow^1 G_1)$.*

Here, agent A_2 needs to compromise, because for every joint plan where he has to work less than he would if he were alone in the world, the other agent would prefer to work alone, and thus will not agree to the joint plan. However, since it is still a dummy compromise encounter, he still does not have to accept the best deal in the sense of the product maximization condition, and he can force A_1 into a deal that is better for him (A_2

Definition 15 *Let $\langle s, (G_1, G_2)\rangle$ be a dummy cooperative/compromise encounter. Let A_1 be the cooperating agent, and let A_2 be the compromising agent. We will say that $\langle s, (G_1, G_2)\rangle$ is a* regular dummy cooperative/compromise encounter *if $\forall J \in NS$ such that $c_2(J) < c(s \longrightarrow^2 G_2)$, it is true that $c_1(J) \geq c(s \longrightarrow^1 G_1)$ and $\exists J \in NS$ such that $c_2(J) = c(s \longrightarrow^2 G_2)$ and $c_1(J) \leq c(s \longrightarrow^1 G_1)$.*

We now give examples of each of the dummy cooperative/compromise encounters described above. Recall our example of a dummy cooperative/compromise encounter from Figure 6. First, note that there is no joint plan that has a role which costs less than 2. So this encounter cannot be a full dummy cooperative/compromise encounter, because A_2 can never do better in a joint plan than he would working alone.

Now consider the case where $w_1 = 5$, $w_2 = 5$. This is a regular dummy cooperative/compromise encounter. Agent A_2 will offer the joint plan in the picture, where he will take role 2 with a probability of 1 and A_1 will take role 1 with a probability of 1. A_1 will agree to this offer, since this gives him a utility of 1 instead of the 0 he would achieve alone. Agent 2 will not have to compromise.

What about the case where $w_1 = 3.5$, $w_2 = 3.5$? With these worths, this is a semi-dummy cooperative/compromise encounter. A_2 must compromise, but has a better deal than the one he will get from the standard negotiation protocol. Let us choose $\varepsilon > 0$. A_2 can offer the joint plan in the picture, where he takes role 2 with a probability of $0.75 - \varepsilon$ and role 1 with a probability of $0.25 + \varepsilon$. A_1 will agree to this offer since it gives him expected work of $3.5 - 2\varepsilon$, and thus an expected utility of 2ε. A_2 has expected work of $2.5 + 2\varepsilon$, and thus expected utility of $1 - 2\varepsilon$, which is better than the 0.5 he would get with the best product maximization plan.

Now consider the encounter in Figure 7. Here $w_1 = w_2 = 10$. Agent A_1 can achieve his task alone, with a cost of 12 pickup/putdown operations, which he cannot afford. Agent A_2 can achieve his goal alone with 4 pickup/putdown operations. The best joint plan has one role that costs 2, and another role that costs 8. The best deal in the sense of the product maximization condition is the one where each agent takes each role with a probability of 0.5. This gives each of the agents an expected utility of 5.

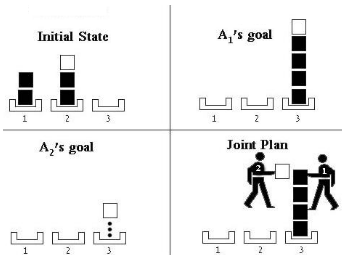

Figure 7. A Full Dummy Cooperative/Compromise Encounter

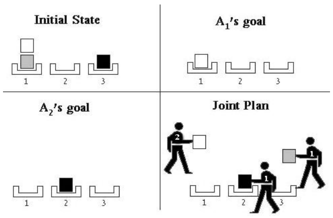

Figure 8. An Example of a True Cooperative/Compromise Encounter

For A_1 this is better than the 0 utility he would get alone. For A_2 this is less than the 6 he would get alone. However, we can see that this interaction answers all the conditions in the definition of a full dummy cooperative/compromise encounter. Agent A_2 can refuse this deal, and offer the deal where he takes the role that costs 2 with a probability of 1 and A_1 takes the role that costs 8 with a probability of 1. Agent A_1 will agree to this deal, since he cannot force A_2 into a better deal (if he will work alone he will do A_2's task on the way). So A_1 will have a utility of 2, and A_2 will have a utility of 8. Now, this cannot be considered a compromise encounter at all—both agents would achieve less if they were alone in the world! This is the reason that this is called a full dummy cooperative/compromise encounter.

2.3.2 Other Types of Cooperative/Compromise Encounters

So far we have dealt with dummy cooperative/compromise encounters. Let us now deal with the other types of cooperative/compromise encounters, the ones where the cooperative agent can force the compromising agent into the compromise.

Definition 16 *Let $\langle s, (G_1, G_2)\rangle$ be a non-symmetric cooperative/compromise encounter in a two-agent SOD; A_1 is the cooperating agent, A_2 is the compromising agent. $\langle s, (G_1, G_2)\rangle$ is said to be a* true cooperative/compromise encounter *if: $\exists g \in G_1$ such that g is achievable alone by A_1 and $\exists J \in$P-NS such that $c_2(J) \leq c(g \longrightarrow^2 G_1 \cap G_2)$. If there is no such $g \in G_1$, $\langle s, (G_1, G_2)\rangle$ is said to be a* true cooperative/compromise encounter *if $\exists J \in$P-NS such that $\forall J' \in NS$, $c_2(J') < c_2(J) \Rightarrow c_1(J') > c(s \longrightarrow^1 G_1)$.*

A full compromise can thus result from two reasons: 1) the cooperating agent can force the compromising agent into the deal, by threatening to work alone; 2) the encounter's nature causes there to exist an optimal joint plan in the sense of the product maximization condition, and there is no other joint plan that is better for the compromising agent that the cooperating agent would accept (i.e., it would be better for him to work alone). Figure 8 shows an example of a true cooperative/compromise encounter.

Here, if A_1 were alone in the world, he would have to use 6 pickup/putdown operations in order to achieve his goal. If A_2 were alone in the world, he would have to use only 2 pickup/putdown

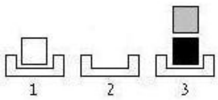

Figure 9. The World After A_1 Works Alone

operations. The best joint plan that satisfies the product maximization condition has one role that costs 2, and one role that costs 4, and each agent will take each role with a probability of 0.5. This gives each agent expected work of 3. Agent A_1 can force A_2 into this joint plan—he can threaten to work alone. Figure 9 shows a possible encounter after A_1 works alone. Now, A_2 would need 4 pickup/putdown operations in order to achieve his goal. Therefore, he would rather use the joint plan that A_1 offered.

3 CONCLUSIONS

We have shown that by using the "threat of passivity," a negotiating agent can achieve better deals than he would achieve using the prior negotiation model of [9]. We formally defined the interactions in which this threat is useful. Furthermore, we showed that adding this threat can change the basic nature of the interactions. For example, a compromise encounter can turn into a cooperative one, where the other agent is suddenly welcomed in the world.

ACKNOWLEDGEMENTS

The first author of this paper gratefully acknowledges receipt of a student travel award from ECCAI.

REFERENCES

[1] Colin Camerer and Richard H. Thaler, 'Anomalies: Ultimatums, dictators and manners', *Journal of Economic Perspectives*, **9**(2), 209–219, (Spring 1995).

[2] Cosmin Carabelea, 'Adaptive agents in argumentation-based negotiation', in *Multi-Agent-Systems and Applications*, pp. 180–187, (2001).

[3] S. Kraus and D. Lehmann, 'Designing and building a negotiating automated agent', *Computational Intelligence*, **11**, 132–171, (1995).

[4] S. Kraus, M. Nirkhe, and K. P. Sycara, 'Reaching agreements through argumentation: A logical model (preliminary report)', in *Proceedings of the 12th International Workshop on Distributed Artificial Intelligence*, pp. 233–247, Hidden Valley, Pennsylvania, (1993).

[5] Ioannis Lionis, *Negotiation with Concurrent Dialogues in Multi-Agent Systems*, Master's thesis, University of London, September 2002.

[6] M. Palatnik and J. S. Rosenschein, 'Long term constraints in multiagent negotiation', in *The Thirteenth International Distributed Artificial Intelligence Workshop*, pp. 265–279, Seattle, Washington, (July 1994).

[7] Iyad Rahwan, Peter McBurney, and Liz Sonenberg, 'Towards a theory of negotiation strategy (a preliminary report)', in *Proceedings of the 5th Workshop on Game Theoretic and Decision Theoretic Agents (GTDT-2003)*, eds., S. Parsons and P. Gmytrasiewicz, pp. 73–80, (2003).

[8] J. S. Rosenschein and M. R. Genesereth, 'Deals among rational agents', in *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pp. 91–99, Los Angeles, California, (August 1985).

[9] J. S. Rosenschein and G. Zlotkin, *Rules of Encounter*, MIT Press, Cambridge, Massachusetts, 1994.

[10] C. Sierra, N. R. Jennings, P. Noriega, and S. Parsons, 'A framework for argumentation-based negotiation', in *Proc. ATAL-97*, eds., M. P. Singh, A. Rao, and M. J. Wooldridge, number 1365 in Lecture Notes in Computer Science, 177–192, Springer-Verlag, (1998).

[11] R. Yagodnick and J. S. Rosenschein, 'Lies in multiagent subadditive task oriented domains', in *The International Workshop on Multi-Agent Systems*, Cambridge, Massachusetts, (October 1998).

[12] G. Zlotkin and J. S. Rosenschein, 'Mechanisms for automated negotiation in state oriented domains', *Journal of Artificial Intelligence Research*, **5**, 163–238, (October 1996).