Multiagent Negotiation Under Time Constraints*

Sarit Kraus†
Department of Mathematics and Computer Science
Bar Ilan University, Ramat Gan, 52900 Israel
sarit@bimacs.cs.biu.ac.il

Jonathan Wilkenfeld
Dept. of Government and Politics
University of Maryland
College Park, MD 20742
jwilkenf@bss2.umd.edu

Gilad Zlotkin‡
Center of Coordination Science
Sloan School of Management
Massachusetts Institute of Technology
1 Amherst St., E40-179
Cambridge, MA 02139
Gilad@mit.edu

Note: This article appears in *Artificial Intelligence journal* 75(2):297–345, 1995.

*This material is based upon work supported by the National Science Foundation under Grant No. IRI-9123460. Some material in this paper appeared in preliminary form in [26, 27]. We thank Karen Lochbaum for her comments.

†Also affiliated with the Institute for Advanced Computer Studies, University of Maryland, College Park.

‡This research was done while the author was at the Computer Science Department in the Hebrew University and was supported by Leibniz Center for Research in Computer Science.

Abstract

Research in distributed artificial intelligence (DAI) is concerned with how automated agents can be designed to interact effectively. Negotiation is proposed as a means for agents to communicate and compromise to reach mutually beneficial agreements. The paper examines the problems of resource allocation and task distribution among autonomous agents which can benefit from sharing a common resource or distributing a set of common tasks. We propose a strategic model of negotiation that takes the passage of time during the negotiation process itself into account. A distributed negotiation mechanism is introduced that is simple, efficient, stable, and flexible in various situations. The model considers situations characterized by complete as well as incomplete information, and ones in which some agents lose over time while others gain over time. Using this negotiation mechanism autonomous agents have simple and stable negotiation strategies that result in efficient agreements without delays even when there are dynamic changes in the environment.

Contents

1 Introduction

Research in distributed artificial intelligence (DAI) is concerned with how automated agents can be designed to interact effectively. One important capability that could aid inter-agent cooperation is negotiation; agents could be built that are able to communicate their respective desires and compromise to reach mutually beneficial agreements.

One of the presumed difficulties in using negotiation as a way of reaching mutual benefit is that negotiation is a costly and time-consuming process and, consequently, it may increase the overhead of coordination (see [1]). In the presence of time constraints, planning and negotiation time should be taken into consideration. The negotiation may be either about job sharing or resource allocation. In both cases we want to prevent the agents from spending too much time on negotiation and therefore not keeping to their timetables for satisfying their goals.

Research in DAI is divided into two basic classes: *Distributed Problem Solving (DPS)* and *Multi-Agent Systems (MA)* [1] (see discussion of previous work in DAI in Section 1.5 below). Research in DPS considers how the work involved in solving a particular problem can be divided among a number of modules or "nodes." The modules in a DPS system are centrally designed to improve performance, stability, modularity, and/or reliability. They include the development of cooperation mechanisms designed to find a solution to a given problem.

Research in MA is concerned with coordinating intelligent behavior among a collection of autonomous (possibly heterogeneous) intelligent (possibly pre-existing) agents. In MA, there is no global control, no globally consistent knowledge, and no globally shared goals or success criteria. There is a possibility for *real competition* among the agents.

These classes are actually the two extreme poles in the DAI research spectrum. Our research falls "closer" to the MA pole since it deals with interactions among *self-motivated*, *rational* and *autonomous* agents. However, we also deal with the possibility that the agents may share a common goal, although even in such situations, the agents are self-motivated and act only according to their interests. We assume that each agent has its own utility function, and that rational behavior involves maximizing expected utility.

We examine the problems of resource allocation and task distribution among autonomous agents. In some domains, agents, due to limited resources *must* share common resource (e.g., roads, bridges, clean air). In other domains, when resources are unlimited, agents may still mutually benefit from sharing a common resource since resources may be expensive (e.g., printers, satellites), or from distributing a set of common tasks. Both problems (resource sharing and task distribution) are symmetrical. In the resource sharing problem there is competition for a valuable resource, with each agent seeking a larger share of the resource. In the task distribution problem where agents have a common goal, several tasks need to be performed to fulfill the goal. Each agent would like the common goal to be achieved with the least amount of effort on its part. This cooperative case also has a competitive element. Each agent wants to perform a smaller part of the job (task).

In this paper we suggest *a strategic model of negotiation that takes the passage of time during the negotiation process itself into consideration.* Changes in the agents' preferences over time will change their strategies in the negotiation and, as a result, the agreements they are willing to reach. This model will show that delays in reaching agreements can be avoided. We will examine the following possible situations where the strategic model is applicable:

1. Two agents that lose over time need to share a common resource. Each agent knows all relevant information about the other agent. They have no alternative, but to continue the negotiations until an agreement is reached (Section 3).

2. Two agents that lose over time need to cooperate to satisfy a common goal. Each agent knows all relevant information about the other agent. The agents can unilaterally leave the negotiations (Section 4).

3. Two agents need to share a resource. One of the agents already has access to the resource and is using it during the negotiation process. It is gaining over time. The other agent is waiting to use the resource and loses over time. Both agents have full information and can unilaterally leave the negotiations (Section 5).

4. Similar to case 3, but the agents do not have complete information about each other (Section 6).

5. Several agents need to cooperate to satisfy a common goal. All of them are losing over time, have full information about each other and can unilaterally leave the negotiations (Section 7).

1.1 The Resource Allocation Problem

A set of agents shares a joint resource. The joint resource can only be used by one agent at a time. Agreement is sought so that all the agents will be able to use the resource. An agreement is a schedule that divides the usage of the resource among the agents.[1]

Examples of joint resources are: communication lines, printers, disks, bridges, road junctions, fresh water, clean air, etc. (Other work in the DAI community dealing with the resource allocation problem includes, for example, [5, 30] which present a multistage negotiation protocol that is useful for cooperatively resolving resource allocation conflicts arising in distributed networks of semi-autonomous problem solving nodes. [32] addresses tradeoff in resource allocation and real-time performance, and develops a mechanism for resource allocation based on the criticality of tasks; [23] proposes resource allocation using specialist "sponsor" agents; and [4] proposes resource allocation via resource pricing).

A communications satellite is a good example of a shared resource, due to the high cost of its launching and maintenance. In many cases the only way a company can get access to a communications satellite is by sharing one with other companies. Even competing companies may find it mutually beneficial to participate in such a joint project.

Sharing a common resource requires a coordination mechanism that will manage the usage of the resource. Discussion about the coordination mechanism will begin (and may even conclude) before discussion of other technical aspects of the joint project.

A coordination mechanism can be a static division of frequencies or time slots. On the other hand, it can be an on-line negotiation mechanism that dynamically resolves local conflicts over the usage of the common resource. Those are the two extreme poles of the coordination mechanism spectrum. On this spectrum there are also coordination mechanisms that generate agreements on long term (an hour, a day, ...) global schedules.

This paper addresses the kinds of attributes a coordination mechanism should have and presents a negotiation mechanism that satisfies those attributes. An important attribute is *efficiency*. The coordination mechanism should result in an efficient joint usage of the common resource. Efficiency implies other attributes such as simplicity and instantaneously (i.e., local conflict should be resolved without delay).

As in the case of the communications satellite, common resources are shared by different companies with possibly different and even conflicting goals. Therefore, to ensure efficiency, the mechanism should also be stable and symmetric. This paper formally defines those attributes (Section 1.3) and presents symmetric, stable and simple on-line coordination mechanisms that resolve local conflicts without delay and result in an efficient joint usage

[1]Our model is also applicable in the case where the resource itself can actually be divided between the agents. This case does not differ significantly from the case where only the resource usage time can be divided.

of the resource. There is some cost associated with the time that elapses between the time that the resource is needed by an agent and the time the agent actually gains access to the resource. This cost depends on the internal state of the agent. For example, its task load, its disk space, etc.

1.2 The Task Distribution Problem

A set of autonomous agents has a common goal it wants to satisfy as soon as possible. In order to satisfy any goal, costly actions must be taken and an agent cannot satisfy the goal without reaching an agreement with the other agents. Each of the agents wants to minimize its costs, i.e., prefers to do as little as possible. We note that even though the agents have the same goal (under our simplified assumptions), there is actually a conflict of interests. The agents try to reach an agreement over the division of labor. We assume that each step of the negotiation takes time, and the agents have preferences for reaching agreements in different time periods (research on the task distribution problem in the area of Distributed Problem Solving systems includes, for example, Davis and Smith's work on the Contract Net [54], Cammarata et al.'s work on strategies of cooperation that are needed for groups to solve shared tasks effectively in the context of collision avoidance in air traffic [2], Lesser and Erman's model of a distributed interpretation system that is able to function effectively even though processing nodes have inconsistent and incomplete information [31], and Carver et al.'s work on agents with sophisticated models that support complex and dynamic interactions between the agents [3]).

An example of task distribution is the "delivery domain" [60, 64, 49, 15]. A group of delivery companies can reduce their overall and individual delivery costs by coordinating their deliveries. Each delivery requirement is a single task. Delivery coordination is actually the exchanging of tasks. One company, for example, that needs to make a delivery from A to B and a delivery from C to D can execute other deliveries from A to B with no extra cost. Therefore, it may agree to exchange its C-to-D delivery with another A-to-B delivery. The mechanisms presented in this paper allow multiple delivery companies to reach an efficient agreement on task distribution without delay that will be mutually beneficial.

1.3 Criteria for Evaluation of Negotiation Protocols

In a multiagent competition situation there is a need to define a mechanism (a protocol) that allows agents to resolve their conflicts and to reach a cooperative agreement. Those mechanisms are usually called *negotiation protocols*.

Given a multi-agent domain, we are interested in investigating both the negotiation protocols that are available to the agents, and also the agent's behavior (negotiation strategy) that is suitable for a given protocol. We will present the optimal strategy an agent should follow in a given protocol, and show that in the design of agents, there should be no reason to adopt any other strategy.

What are the conditions that a *Negotiation Protocol* should satisfy (for any specific distributed multi-agent domain), such that it should be accepted by all the designers of agents (for that specific domain)?

- **Distributed** — The decision making process should be distributed. There should be no central unit or agent that is managing the process.

- **Instantaneously** — Conflict should be resolved without delay.

- **Efficiency** — The outcome of the negotiations (i.e., the agreements) should be efficient:

  - Conflict should be avoided when possible and the mechanism should allow the agents to reach *Pareto optimal* agreements with high probability. An agreement is Pareto optimal if there is no other agreement that dominates it, i.e., there is no other deal that is better for some of the agents and not worse for the others.

  - In the resource allocation problem, the resource is not in use only when there is no agent in the group that currently needs the resource (there are no deadlocks).

- **Simplicity** — The negotiation process itself should be simple and efficient. It should be short and consume only a reasonable amount of communication and computation resources.

- **Symmetric** — The coordination mechanism should not treat agents differently because of non-relevant attributes. In the situations that we consider, the agents' utility functions and their role in the encounter are the relevant attributes. All other attributes, like an agent's color, name or manufacture are not relevant. That is, Symmetric implies that given a specific situation, the replacement of an agent with another which is identical with respect to the above attributes, will not change the outcome of the negotiation.

- **Stability** — There should be a distinguishable (Nash or even subgame-perfect) equilibrium point to the negotiation protocol (considered as a game).[2] Given a specific situation, we would like to be able to find simple strategies that we could recommend to **all** agent designers to build into their agents. No designer will benefit by building agents that use any other strategy. The equilibrium point should not violate the efficiency condition, i.e., the negotiation should result in a Pareto-optimal agreement. Being a "simple strategy" means that it is feasible to build it into an automated agent. A "simple strategy" also presumes that an agent will be able to compute the strategy in a reasonable amount of time.

- **Satisfiability or accessibility**— In the resource allocation case we would like an agent that needs the resource to eventually have access to the resource (there is no starvation). In the task distribution case, we would like the task to eventually be performed.

1.4 Related Work in DAI

The study of multi-agent interaction has been receiving increasing attention within artificial intelligence (AI). This is a direct outgrowth of the serious consideration currently being given to agents operating in challenging, real-world environments. For many years, highly restricted domains were considered sufficient for AI research purposes, and agents such as Shakey [14] could be designed and built for operation in simplified, restricted environments.

The research on agent architectures and on planning typically made several standard assumptions, including the existence of a static domain, the lack of deadlines, and the existence of a single agent, i.e., *our* agent. Once researchers began, for a variety of reasons, to move into realistic domains, these assumptions had to be quickly discarded. The research in planning and agent architectures of the last decade has been focused precisely on the transformation of single-agent, atemporal, static theories into multi-agent, temporal, dynamically capable ones.

Researchers on agent interaction differ over the basic assumption of the degree of control that the designer has over individual agents and over their social environment (i.e., interaction mechanisms). Therefore, we can make a two dimensional classification. On the first dimension we have the degree of control over the *social layer* of the agents. It ranges from a

[2]For additional discussion of the concepts of Nash and subgame perfect equilibrium, and Pareto optimality, see Sections 1.5 and 2.1 below.

	Structured	Unstructured
DPS	Moses, Shoham & Tennenholtz [38, 53] Davis & Smith [54] Malone [37] Lesser [33, 6] Durfee [11, 12]	
MA	Zlotkin &Rosenschein [64, 45, 61] Wellman [60], Ephrati & Rosenschein [13] Kraus, Wilkenfeld & Zlotkin	Sycara [57, 58], Kraus & Lehmann [24, 25] Grosz [34, 19], Gasser [17]

Figure 1: The rows indicates Degree of control on the "social Layer". The columns indicates degree of control on other agents in the domain.

highly structured interaction mechanism to a totally unstructured interaction (See Figure 1).

On the second dimension we have the degree of control that a designer has over individual agents. It ranges from the case where a single designer is able to control (or even explicitly design) each individual agent in the domain (those systems are known as Distributed Problem Solving (DPS)), to the case where there are multiple designers and each is able to design only its agent and has no control over the internal design of other agents in the domain (those systems are known as Multi-Agent system (MA)).

The second dimension is also tightly coupled to the issue of agents' incentives. When agents are assumed to be centrally designed they are also assumed to have a common general goal. In such cases each agent tries to maximize some system global utility. When agents are designed by different designers agents, they are usually assumed to have individual motivation to achieve their own goal and to maximize their own utility.

In the upper left corner of our two dimensional matrix (i.e., designer can fully control each individual agent and also the interaction environment) we find the work of Shoham, Tennenholtz and Moses on Social Laws [53, 38] which shows that "pre-compiled" highly structured "social laws" are able to coordinate agent activity and to restrict on-line conflict. Agents are assumed to follow the "social laws" since they were designed to and not because they individually benefit from the "social laws." The same approach could have been applied in an MA system in the case were the social laws are "stable", i.e., it is in each agent's individual interest to follow the law. In our research we assume that the agents are individually motivated and therefore the issue of stability plays an important role in the design of the interaction mechanism.

Even in DPS systems it may be useful to incorporate pure competition among the agents. Davis and Smith's work on the Contract Net [54] introduced a form of simple negotiation among cooperative agents, with one agent announcing the availability of tasks and awarding them to other bidding agents. Malone refined this technique considerably by overlaying it with a more sophisticated economic model [37], proving optimality under certain conditions. In the general contract net approach and also in the economic oriented refinements the main underlying assumption is that agents are "benevolent" and are motivated to help each other. Such an assumption is not feasible when agents are self motivated.

Another more experimentally based approach for inter-agent collaboration in DPS is presented in the on-going research of Lesser, Durfee and their colleagues using the "Functionally Accurate, Cooperative, (FA/C)" paradigm. For example the "Partial Global Planning" architecture has been implemented and evaluated in the vehicle monitoring domain (DVMT) ([11, 12, 7]). The agents iteratively exchange tentative partial solutions to construct global solutions. Multi-agent planning in the communication network domain was treated as a distributed constraint on satisfaction and was implemented by using a multi-stage negotiation [6]. The multi-stage negotiation provides each agent sufficient information to enable it to make local decisions that are globally correct. These researchers approach the issue of global efficiency and performance more directly in real world working systems, while we are analyzing the use of formal tools and general mechanisms in more idealized domains.

We are unfamiliar with work that belongs to the upper right corner where the agents are centrally designed, but use unstructured communication protocols. This is, of course, due to the fact that since structured communications protocols usually provide more efficient cooperation, and if the designer has control over all the agents, it can incorporate a structured communication protocol in the agents to make the DPS system more efficient.

In the lower right corner the designers have control only over their agents (i.e., no control over other agents or the interaction mechanism). This is usually the case in domains where humans are interacting with each other and with autonomous agents. For example, in the case of labor negotiation, Sycara [57, 58] presented a model of negotiation that combines case-based reasoning and optimization of the multi-attribute utilities. In her work agents try to influence the goals and intentions of their opponents. In [25, 24] Kraus and Lehmann developed an automated Diplomacy player that negotiates and plays well in actual games against human players. Researchers in discourse understanding (e.g., [34, 19]) develop formal models to support communication in human/machine interaction. There, models of individual and shared plans are used for understanding non-structured communication.

Gasser ([17]) focuses on the social aspects of agent knowledge and action in multi-agent systems ("communities of programs"). As in real-world societies, social mechanisms can dynamically emerge. Communities of programs can generate, modify, and codify their own local languages of interaction. Gasser's approach may be most effective when agents are interacting in unstructured domains, or in domains where their structure is continuously changing. In our research, we choose to pre-design the social layer of multi-agent systems by creating a structured interaction mechanism, i.e., a model of alternating offers.

In most of the unstructured negotiation scenarios there is no guarantee that agreement will be reached, and the negotiation may take a long time. In the current work the negotiation always ends at the latest in the second stage of the negotiation and if there is complete information, agreement is guaranteed.

The work we describe in the present paper, resides in the lower left corner. It assumes that there is full control over the agent interaction mechanism by bounding the agents to highly structured public behavior (like negotiation protocols, voting procedures, bidding mechanisms, etc.). However, as we mentioned in the introduction, our work is concerned with problems in developing agents in *Multi-Agent Systems*. That is, there is no control over the other agent's private behavior. This gap is bridged by carefully adjusting the interaction mechanism such that it will be stable. Using a stable mechanism, it is to the benefit of each individual agent (that wishes to maximize its own private utility) to adopt a given private behavior. When those private behaviors (strategies) are in equilibrium, then the designers of the interaction protocols can assume that the individual agents will be designed to have those private behaviors even though the protocols' designers have no explicit control.

Ephrati and Rosenschein ([13]) used the Clarke Tax voting procedure as a consensus mechanism. The mechanism assumes an explicit utility transferability (i.e., a kind of monetary system). In the problem of task distribution and resource allocation there is no explicit way to transfer utility. There is an *implicit* ways to transfer utility, e.g., by executing one of your tasks I may transfer some utility I could have been getting to you. However, this implicit utility transfer is not sufficient for the implementation of the Clark Tax procedure. The Clark Tax mechanism assumes that agents are able to transfer utility out of the system (the taxes that are being paid by the agents). The utility that is transferred out of the system is actually wasted and reduces the efficiency of the consensus that is reached. This is the price that needs to be paid to ensure stability. In the paper we introduce a negotiation mechanism that provides both efficiency and stability.

Zlotkin and Rosenschein [64, 45, 61] analyze the relationship between the attributes of the domain in which the agents are operating and the availability of interaction mechanisms to satisfy the efficiency, simplicity, symmetry, and stability conditions. They have classified interaction domains as task oriented domains, state oriented domains and worth oriented domains. In all of the above domains time plays no explicit role in the agent's utility functions. It may be appropriate when negotiation time can be neglected relative to plan execution time. However, in highly dynamic systems, negotiation time plays an important role in the evaluation of the performance of the system and cannot be neglected. The approach presented in this paper focuses precisely on this kind of domain and provides coordination mechanisms that ensure efficient agreements with no delay. Within the MA/structured group of researchers it is the first attempt to treat the temporal aspect of negotiation explicitly.

In this paper we consider the problem where agreements involve all the agents. Multi-agent's negoiation mechanisms, in situations in which agents are free to form any coalition that includes some of the agents while exluding others, are discussed in [52, 62, 22].

Wellman ([60]) uses a market oriented approach for inter-agent coordination mechanism design ("market oriented programming"). When the agent interaction can be reduced to a simple consumer-producer relation a market pricing mechanism can be used to ensure efficiency and stability. However, not all inter-agent interaction can be mapped to the consumer-producer paradigm. For example, in the case of a common resource (i.e., one of the encounters that are considered in the present approach), the whole community of agents are consumers.

To summarize, our work is characterized by providing a formal strategic model of negotiation that takes the passage of time during the negotiation process itself into account. It can be used for both resource allocation and task distribution, without side payments. The only assumptions that we made is on the negotiation protocol, which is a protocol of alternating offers which we describe in details in Section 3. However, we don't make any assumptions about the offers the agents make during the negotiation as is the case in some other work (e.g., [63]). In particular, the agents are not bounded to any previous offers that have been made. Nevertheless, the negotiation ends with no delay.

1.5 Related Work in Economics and Game Theory

There are two main approaches to the development of theorems relating to the negotiation process. The first is informal theories which attempt to identify possible strategies for a negotiator and to assist a negotiator in achieving optimal results (see [16, 10, 21]). The other approach is the formal theory of bargaining originating with the work of John Nash ([39] [40]), who attempted to construct formal models of negotiation environments and to prove different theorems about the best strategies a negotiator can follow under different circumstances.

This formal game theory approach provides clear analyses of various situations and precise results concerning the strategy a negotiator should choose. However, it requires making restrictive assumptions that are unacceptable to the first group.

Following Genesereth, Ginsberg, Rosenschein and Doyle, [18, 44, 8, 9], we propose the use of game-theoretic techniques for Artificial Intelligence purposes. We propose to develop a strategic model of negotiation that can serve as the basis for building efficient automated negotiators. The formal game theory approach is also divided into two central sub-approaches concerning the bargaining problem (see [20]). The first is the strategic approach. The agents' negotiating maneuvers are moves in a noncooperative game and the rationality assumption is expressed by investigation of the *Nash Equilibrium*.[3]

The second approach is the axiomatic method. It makes assumptions about the solution of a negotiation situation without specifying the bargaining process itself (the literature on the axiomatic approach to bargaining is surveyed by Roth [46], and can also be found in [35] with a general introduction to game theory.)

Since we intend to use our theoretical work as a basis for the development of automated negotiators, we have adopted the strategic approach. Rubinstein [47] and Ståhl [56] developed models of alternating offers, which take time into consideration. Shaked and Sutton [51] extended these works by developing models in which an agent can opt out of the game. Those works are closely related to our desired models (see [41] for a detailed review of the bargaining game of alternating offers). Nevertheless, several important modifications are needed. These mainly concern the way time influences the preferences of the agents, the possibility that both agents can opt out, and the preferences of the agents over opting out. Only the results in Section 3.2 are based on Rubinstein's previous work; all the other results are ours.

2 Two Fully Informed Agents

In the next three Sections we consider the case where two fully informed agents negotiate to reach agreement on resource allocation or on task distribution.

These situations are characterized by the following assumptions.

1. **Bilateral Negotiation** — Even if there are several agents in the environment, the initial assumption is that in a given period of time no more than two agents need the same

[3] A pair of strategies (σ, τ) is a Nash Equilibrium if, given τ, no strategy of Agent 1 results in an outcome that Agent 1 prefers to the outcome generated by (σ, τ) and similarly for Agent 2 given σ.

resource (we will relax this assumption in Section 7). When there is an overlap between the time segments in which two agents need the resource, these agents will be involved in a negotiation process.

2. **Full Information** — Each agent knows all relevant information including the other agent's utilities for the different outcomes over time (we will relax this assumption in Section 6).

3. **Rationality** —The agents are rational; they try to maximize their utilities and behave according to their preferences.

4. **Commitments are Kept** — If an agreement is reached both sides will honor it.

5. **No Long Term Commitments** — Each negotiation stands alone. An agent cannot commit itself to any future activity other than the agreed-upon schedule.

6. **Resource Division Possibilities** — We assume that any division of the resource is possible (we will relax this assumption in Section 4 through the end of the paper).

7. **No Other Options** — The agents have no alternative, but to continue the negotiations until an agreement is reached (we will relax this assumption in Section 4 through the end of the paper).

8. **Common Knowledge** Assumptions (1)-(7) are common knowledge.

2.1 Strategies and Equilibrium

Our strategic model of negotiation is a model of Alternating Offers.[4]

How will a rational agent choose its negotiation strategy in such an environment? A useful notion is the Nash Equilibrium [40, 35] . If there is a unique equilibrium, and if it is known that an agent is designed to use this strategy, no agent will prefer to use a strategy other than this one.

However, the use of Nash equilibrium is not an effective way of analyzing the outcomes of the models of Alternating Offers since there may be some points in the negotiation where one or more agents prefer to diverge from their Nash equilibrium strategies. Nash equilibrium strategies may be in equilibrium only in the beginning of the negotiation, but may be unstable in intermediate stages.

[4]See [41] for a detailed review of the bargaining game of Alternating Offers.

Nash equilibrium puts few restrictions on the outcome and also yields too many equilibrium points (see ([47] for the proof).

Therefore, we will use the stronger notion of **(subgame) perfect equilibrium** (P.E.) (see [50, 47]) which requires that the agents' strategies induce an equilibrium at any stage of the negotiation, i.e., in each stage of the negotiation, assuming that an agent follows the P.E. strategy, the other agent has no strategy better than to follow its own P.E. strategy. Subgame perfect equilibrium is essentially a backward induction argument, using the rationality of agents at each stage of the game to decide what a good choice is and then rolling backward [59]. So, if there is a (unique) perfect equilibrium, and if it is known that an agent is designed to use this strategy, no agent will prefer to use a strategy other than this one in each stage of the negotiations.

We will consider different variations of this model. In the first case, we assume that the agents are bound to an agreement, i.e., the negotiation process can end only by reaching an agreement. Otherwise, the agents will continue to negotiate forever. In the second case, the agents are able to opt out at any stage of the negotiation.

3 The Bounding Negotiations Mechanism

When agents are bounded to an agreement, negotiation may continue forever. The driving force for reaching an agreement in a reasonable amount of time is the agents' attitudes toward negotiation time. We assume that negotiation time is expensive and taken into consideration by each agent. Even though the negotiation can continue indefinitely we will show that agreement will be reached without any delay.

We utilize modified definitions from [41]. We assume that there is a set of agents $\mathcal{A} = \{1, 2\}$. We present a formal definition of an agreement.

Definition 1 Agreement:
An agreement is an ordered pair (s_1, s_2), in which s_i is agent i's portion of the resource. The set of possible agreements is

$$S = \{(s_1, s_2) \in I\!R^2 : s_1 + s_2 = 1 \text{ and } s_i \geq 0, \text{ for } i = 1, 2\}$$

Each agent has a preference over the set of possible agreements S. Our sole assumption is that an agent prefers an agreement that gives it a larger portion of the resource over an agreement that gives it less.

Negotiation is a process that may include several iterations and may even continue forever. We assume that agents can take actions only at certain times in the set $\mathcal{T} = \{0, 1, 2...\}$ that are fixed in advanced. In each period $t \in \mathcal{T}$ one agent, say i, proposes an agreement from S, and the other agent (j) either accepts the offer (**Yes**) or rejects it (**No**). If the offer is accepted (j says **Yes**), then the negotiation ends with implementation of the agreement (i.e., the resource is used according to the agreement). After a rejection, the rejecting agent then has to make a counteroffer and so on. There are no rules which bind the agents to any specific strategy. In particular, the agents are not bound to any previous offers that have been made. The mechanism only provides a framework for the negotiation process and specifies the termination condition, but there is no limit on the number of periods. An agent's negotiation strategy in general is any function from the history of the negotiations to its next move.

Definition 2 Negotiation Strategies:
A strategy *is a sequence of functions $f = \{f^t\}_{t=0}^{\infty}$. The domain of the i^{th} element of a strategy is a sequence of offers of length i (all possible histories up to period i) and its range is the set $\{\textbf{Yes,No}\} \cup S$ (its current move).*

That is, if f is a strategy for the first agent to make an offer (agent 1) then $f^0 \in S$ and for t even $f^t : S^t \to S$, and for t odd $f^t : S^{t+1} \to \{\textbf{Yes,No}\}$ (S^t is the set of all sequences of length t of elements in S and **Yes,No** *are defined above). We denote by F the set of all strategies of the agent which starts the bargaining. Similarly, let G be the set of all strategies of the agent which, in the first move, has to respond to the other agent's offer; that is, G is the set of all sequences of functions $g = \{g^t\}_{t=0}^{\infty}$ such that for t even $g^t : S^{t+1} \to \{\textbf{Yes,No}\}$ and for t odd $g^t : S^t \to S$.*

Let $\sigma(f, g)$ be a sequence of offers in which agent 1 starts the bargaining and adopts $f \in F$, and agent 2 adopts $g \in G$. Let $\text{Length}(f, g)$ be the length of $\sigma(f, g)$ (where the length may be infinite). Let $\text{Last}(f, g)$ be the last element of $\sigma(f, g)$ (if there is such an element). We present a formal definition for the outcome of the negotiation.

Definition 3 Outcome of the negotiation:
The outcome function of the negotiation is defined by

$$Outcome(f,g) = \begin{cases} \textbf{Disagreement} & \text{if } Length(f,g) = \infty \\ (Last(f,g), Length(f,g) - 1) & \text{otherwise} \end{cases}$$

Thus, the outcome (s,t) where $s \in S$ is interpreted as the reaching of agreement s in period t and the symbol **Disagreement** indicates a perpetual disagreement.

We note here that by defining an outcome to be either a pair (s,t) or **Disagreement**, we have made a restrictive assumption about the agents' preferences. We assume that agents care only about the nature of the agreement, and the time at which the outcome is reached, and not about the sequence of offers and counteroffers that leads to the agreement, i.e., there is no "decision-regret" (see [42]).

3.1 The Agents' Utility Functions

We also assume that agent $i \in \mathcal{A}$ has a continuous utility function over all possible outcomes: $U^i\colon \{S \times \mathcal{T}\} \cup \{\textbf{Disagreement}\} \to I\!R$. Throughout the paper, when the utility for an agent from one outcome is greater than from another outcome, we will assume that the agent prefers the first outcome over the second (Due to Assumption 3—Rationality, in Section 2).

With the exception of Section 6 which considers situations with incomplete information, none of the results reported in this paper depend on the exact values of the utility functions over the possible outcome. The factor that plays the key role in reaching a specific agreement is the relation among the utility values of the outcomes. That is, whether the utility of an agreement s is greater than the utility s' for a given agent, and not the exact utility values of s and s'. The results can be obtained in systems where the agents have preferences on the possible outcomes and not numerical utility functions. One of the issues related to the application of these results is the computing of these utility functions, or determining the preferences; this will be discussed briefly in Section 9.

3.1.1 Attributes of the Utility Functions

We now present a number of assumptions concerning the utility functions of the agents. This basic set of assumptions will be added to and modified in subsequent sections as we introduce additional conditions such as incomplete information, opting out, time, and multiple agents. Subscripts denote the section to which the specific assumption refers.

The first assumption states that agents prefer any agreement in any given time period over the continuation of the negotiation process indefinitely.

$A0_3$ **Disagreement is the worst outcome:** For every $s \in S, i \in \mathcal{A}$ and $t \in \mathcal{T}$, $U^i((s,t)) > U^i(\textbf{Disagreement})$.

The next two conditions ($A1_3$), ($A2_3$) concern the behavior of the utility function U^i on $S \times \mathcal{T}$, i.e., agreements reached in different time periods. Condition ($A1_3$) requires that among agreements reached in the same period, agent i prefers larger portions of the resource.

$A1_3$ **The Resource is Valuable:** For all $t \in \mathcal{T}, r, s \in S$ and $i \in \mathcal{A}$: $r_i > s_i \Rightarrow U^i((r,t)) > U^i((s,t))$.[5] For agreements that are reached within the same time period, each agent prefers to get a larger portion of the resource.

The next assumption states that time is valuable to both sides.

$A2_3$ **Time is Valuable** For any $t_1, t_2 \in \mathcal{T}, s \in S$ and $i \in \mathcal{A}$ if $t_1 < t_2$, $U^i((s,t_1)) \geq U^i((s,t_2))$.

The next assumption greatly simplifies the behavior of the utility function for agreements. It requires that the difference in utility between (s_1,t_1) and (s_2,t_2) depends only on s_1, s_2 and the differences between t_1 and t_2.

$A3_3$ **Stationarity:** For all $r, s \in S$, $t_1, t_2, \delta \in \mathcal{T}$ and $i \in \mathcal{A}$; $U^i((r,t_1)) \geq U^i((s, t_1+\delta))$ iff $U^i((r,t_2)) \geq U^i((s,t_2+\delta))$

By assumption ($A2_3$) the agents prefer to receive any given share of the resource sooner rather than later. The following assumption imposes the condition that the loss associated with any given amount is an increasing function of that amount.

$A4_3$ **Increasing Loss** For every $i \in \mathcal{A}$, $t \in \mathcal{T}$ and $s \in S$, there exists $s' \in S$ such that $U^i((s,t)) = U^i((s',0))$. Furthermore, for every $s, r, s', r' \in S$ such that $U^i((s,t)) = U^i((s',0))$ and $U^i((r,t)) = U^i((r',0))$ if $s'_i \geq r'_i$ then $s_i - s'_i \geq r_i - r'_i$.

3.1.2 Examples of Utility Functions

We will examine two examples of utility functions which conform to Assumptions $A0_3 - A4_3$.

Time Constant Discount Rates: In the first case, we consider a utility function with a time constant discount rate. That is, every agent i has a fixed discount rate $0 < \delta_i < 1$. If the agents reach an agreement in time period t in which agent i's portion of the resource is s_i, then its utility will be $s_i \delta_i^t$.

[5] For all $s \in S$ and $i \in \mathcal{A}$, s_i is agent i's portion of the resource.

Definition 4 Utility function with time constant discount rate
Let $(s,t) \in S \times \mathcal{T}$ be an outcome of the negotiation, then the $Utility_i((s,t))$ where $i \in \mathcal{A}$ is defined to be $\delta_i^t s_i$, where $0 < \delta_i < 1$, and $Utility_i\{\textbf{Disagreement}\} = -\infty$.[6]

Constant Cost of Delay: The second case is of a utility function with a constant cost due to delay. Here, every agent bears a fixed cost for each period. That is, each agent i has a constant $c_i > 0$, and if the agents reach an agreement in time period t in which agent i's portion of the resource is s_i, then its utility will be $s_i - c_i t$. Formally:

Definition 5 Utility function with a constant cost of delay
Let $(s,t) \in S \times \mathcal{T}$ be an outcome of the negotiation; then the $Utility_i'\{(s,t)\}$ where $i \in \mathcal{A}$ is defined to be $s_i - c_i t$ where $c_i > 0$ and $Utility_i'\{\textbf{Disagreement}\} = -\infty$.

A utility function with a time constant discount rate satisfies all the above conditions, while a utility function with a constant cost due to delay satisfies all but $A4_3$.

3.2 Equilibrium Yields Agreement with No Delay

In [47] it was proved that for continuous utility functions that satisfy axioms $A0_3 - A4_3$ there exists a unique P.E. which results in the successful termination of the negotiation after the first period.[7] This unique solution is characterized by a pair of agreements x^* and y^* that satisfy these conditions: (1) Agent 1 is indifferent between "y^* today" and "x^* tomorrow," and (2) Agent 2 is indifferent between "x^* today" and "y^* tomorrow." When a unique pair of x^* and y^* satisfies this statement, there exists a unique P.E. [47]. The structure of the unique perfect equilibrium is as follows: Agent 1 (2) always suggests x^* (y^*) and agent 2 (1) accepts any offer which is at least as good for it as x^* (y^*).

We will now demonstrate the usage of these results for the two types of utility functions: (1) a constant discount rate; and (2) a constant cost of delay.

3.2.1 Constant Discount Rate

In the case of a utility function with a constant discount rate, as in definition 4, agent 1 has a discount rate of $0 < \delta_1 < 1$, and agent 2 has a discount rate of $0 < \delta_2 < 1$, where

[6] Here, $\delta_i^t s_i$ denotes δ_i to the t-th power times s_i.

[7] Rubinstein's results are actually more general, and consider the case of agents' preferences in addition to their utility functions. We use utility function to be consistent with our approach in section 6.

$Utility_i((s,t)) = \delta_i^t(s_i)$. According to the P.E. strategies, in every period of time, when it is agent 1's turn to make an offer, it will offer agent 2 $(\frac{1-\delta_2}{1-\delta_1\delta_2}, \frac{\delta_2(1-\delta_1)}{1-\delta_1\delta_2})$. When agent 1 receives an offer from agent 2, agent 1 will accept only the offers where its share of the resource is at least $\frac{\delta_1(1-\delta_2)}{1-\delta_1\delta_2}$.

On the other hand, when it is agent 2's turn to make an offer, it will offer $(\frac{\delta_1(1-\delta_2)}{1-\delta_1\delta_2}, \frac{1-\delta_1}{1-\delta_1\delta_2})$. 2 will accept an offer only if its share in it is greater than $\frac{\delta_2(1-\delta_1)}{1-\delta_1\delta_2}$.

The agreement that will be reached in the first period is $(\frac{1-\delta_2}{1-\delta_1\delta_2}, \frac{\delta_2(1-\delta_1)}{1-\delta_1\delta_2})$.

Formally:

Lemma 1 *(***Rubinstein***[47]):*
Suppose agent 1 *starts the negotiations. Let*

$$x^* = (\frac{1-\delta_2}{1-\delta_1\delta_2}, \frac{\delta_2(1-\delta_1)}{1-\delta_1\delta_2}) \quad y^* = (\frac{\delta_1(1-\delta_2)}{1-\delta_1\delta_2}, \frac{1-\delta_1}{1-\delta_1\delta_2})$$

$(\hat{f}, \hat{g})$ *is a subgame perfect equilibrium of the strategic model of Alternating Offers where the agents' utility function is defined in definition 4 iff*
$\hat{f}^t(s^0, ..., s^{t-1}) = x^*$ *for all* $(s^0, ..., s^{t-1}) \in S^t$*, if* t *is even, and*

$$\hat{f}^t(s^0, ..., s^t) = \begin{cases} \textbf{Yes} & \textit{if } \ s_1^t \geq y_1^* \\ \textbf{No} & \textit{if} s_1^t < y_1^* \end{cases}$$

if t *is odd. The strategy* $\hat{g}$ *of agent* 2 *has the same structure; the roles of* x^* *and* y^* *are reversed, the words "odd" and "even" are interchanged, and each subscript agent* 1 *is replaced by* 2*. The outcome is that agent* 1 *proposes* x^* *in the first period (period* 0*), and* 2 *immediately accepts this offer.*

Proof: The proof and additional discussion can be found in [41].

Even though the structure of the strategic model of Alternating Offers allows negotiation to continue indefinitely, in the unique subgame perfect equilibrium it terminates immediately. The exact agreement that will be reached depends mainly on the patience of the agents. An agent whose losses over time are less than its opponent's will get a larger share of the resource. For example, if the agents have a utility function with discount rates then being less patient means having a larger value of δ_i. That is, if δ_1 is smaller, agent 1's share of the resource is smaller, while if δ_2 is smaller, agent 1's share is larger.

In addition, the agent which starts the negotiation has an advantage over the other agent (for example, if both agents have the same rate of delay-δ, then the first one will receive $\frac{1}{1+\delta}$ of the resource and the other will receive $\frac{\delta}{1+\delta}$). A simple way to avoid this asymmetry in the model is the following (see [41]): at the beginning of each period each agent is chosen with probability $\frac{1}{2}$ (independently across periods) to be the one to make the first offer.

3.2.2 Constant Cost of Delay

In the second case, suppose the agents' utility function includes a constant cost of delay as defined in Definition 5. That is, $Utility_i'\{(s,t)\} = s_i - c_i t$ where $c_i > 0$ and where $i \in \mathcal{A}$. Even though this utility function does not satisfy $(A4_3)$, there is a unique P.E. if $c_1 \neq c_2$.

Suppose agent 1 is more patient than agent 2. That is, it loses less over time than agent 2 (i.e., $c_1 < c_2$). In such a situation agent 1 is indifferent between "$(1 - c_1, c_1)$ today" (i.e., $y^* = (1 - c_1, c_1)$) and "(1,0) tomorrow" (i.e., $x^* = (1, 0)$) and agent 2 is indifferent between "(1,0) today" and "$(1-c_1, c_1)$ tomorrow." Therefore, if it is agent 1's turn to make an offer it will always offer $(1, 0)$. Agent 1 will accept any agreement in which its share of the resource is greater or equal to $1 - c_1$. 2 will accept any agreement (including $(1, 0)$), and will offer $(1 - c_1, c_1)$. In the agreement that will be reached agent 1 will use the resource alone. The prediction here is quite surprising. Since agent 1 is more patient than agent 2, it can gain *all* the resource. Agent 2 prefers it over waiting an additional period. If $c_1 > c_2$, that is, agent 2 is more patient than agent 1, $x^* = (c_2, 1 - c_2)$ and $y^* = (0, 1)$. That is, when it is agent 1's turn to make an offer it will offer agent 2 $(c_2, 1 - c_2)$, and will accept any offer. Agent 2 will accept $1 - c_2$ and will offer $(0, 1)$. Here also the results are quite extreme. Agent 1's share will only equal agent 2's delay. Here again, the agent whose turn it is to make the first offer is in a better position than the one which goes second. [8]

These results demonstrate that introducing the time factor into the negotiation process can lead to an efficient negotiation.

4 Unbounded Negotiation Mechanism When Both Agents Lose Over Time

Up to this point we have assumed that the agents have no choice but to continue the negotiation since disagreement is the worst outcome to both sides. Let us consider the case in which the agents have the ability to unilaterally opt out of the negotiation. This can happen when the agents try to satisfy a common goal but the agents have some other goals they

[8] Where $c_1 = c_2$ there are multiple subgame perfect equilibria.

can satisfy (usually with lower priority); a threat to leave the negotiation may influence the outcome in some cases.

In the previous section, we have also assumed that the agents can divide the work/resource between them, in any way that they have agreed upon. Unfortunately, this cannot usually be done. If two agents need to carry blocks, to deliver packages or to build tools, this work can be divided only in a discrete manner and usually in a finite number of possible agreements. From now on we will consider the case of a finite discrete case. That is, the simplifying assumptions 1-5 (Bilateral Negotiation, Full Information, Rationality, Commitments are Kept, No Long Term Commitments) and assumption 8 (Common Knowledge) that are described at the beginning of Section 2 are still valid, while assumptions 6 (Resource Division Possibilities) and 7 (No Other Options) are no longer valid. We assume that there are M units of the work (or resource) that must be divided by two agents.

Example 1 *There are two agents that are responsible for the delivery of electronic newsletters of two different companies. The delivery is done by phone (either by fax machines or electronic mail). The expenses of the agents depend only on the number of phone calls. Therefore, if there is someone who subscribes to both companies' newsletters, the two newsletters may be delivered to it by one of the agents for the price of only one phone call. The agents negotiate over the distribution of the common subscriptions. Each of the agents can opt out of the negotiations and deliver all of its own newsletters by itself.*

We slightly modify the definition of an agreement (Definition 1). The set of possible agreements, S, includes all the pairs $(s_1, s_2) \in I\!N^2$ where $s_1 + s_2 = M$. We also modify the negotiation strategies (Definition 2) such that if agent i receives an offer from its partner it can opt out of the negotiation (**Opt**), in addition to accepting the offer (**Yes**) or rejecting it (**No**).

$\sigma(f, g)$, Length(f, g), Outcome(f, g) and the outcome of the negotiations are defined as in Section 2, but Outcome(f, g), which is the last element of $\sigma(f, g)$, can be either $s \in S$ or **Opt**. Thus the outcome $(\mathbf{Opt}, t)$ is interpreted as one of the agents opting out of the negotiation at period t. We note that the length of the time periods is fixed. The agents' utility functions in this case are over agreements reached at various points in time, and over opting out at various points in time. That is, $U^i\colon \{\{S \cup \{\mathbf{Opt}\}\} \times \mathcal{T}\} \cup \{\mathbf{Disagreement}\} \rightarrow I\!R$.

4.1 Attributes of the Utility Functions

The following set of assumptions, some of which are modifications of the original set presented in Section 3, are necessary to model the negotiation situation in which the parties can choose to opt out.

Condition ($A0_3$) (Disagreement is the worst outcome) of Section 3 is still valid. That is, disagreement is even worse than opting out. Formally we state it as follows:

$A0_4$ **Disagreement is the worst outcome:** For every $s \in S, i \in \mathcal{A}$ and $t \in \mathcal{T}$, $U^i((s,t)) > U^i(\mathbf{Disagreement})$ and $U^i((\mathbf{Opt},t)) > U^i(\mathbf{Disagreement})$.

In this section we deal with the case of task distribution. In such situations, we assume that each agent prefers to do as little as possible. Therefore, condition ($A1_3$) (which was appropriate to the resource allocation case) is modified. We denote the modified condition by $A1_4$. That is ($A1_4$) requires that among agreements reached in the same period, agent i prefers smaller numbers of units s_i.

$A1_4$ **Actions Are Costly:** For all $t \in \mathcal{T}, r, s \in S$ and $i \in \mathcal{A}$: $r_i > s_i \Rightarrow U^i((r,t)) < U^i((s,t))$. For agreements that are reached within the same time period, each agent prefers to perform a smaller portion of the labor.

($A2_3$) (Time is valuable) is still valid. We denote it by ($A2_4$).

We will consider the case of constant delay, in which any agent has a number $c_i > 0$ $i \in \{1,2\}$ that satisfies the following condition.[9]

$A3_4$ **Agreement's Cost Over Time** Each agent $i \in \{1,2\}$ has a number $c_i > 0$ such that: $U^i((s,t_1)) \geq U^i((\bar{s},t_2))$ iff $(s_i + c_i t_1) \leq (\bar{s}_i + c_i t_2)$.

We note that assumption ($A3_4$) does not hold for **Opt**. We also assume that both agents prefer to opt out sooner rather than later. Formally:

$A4_4$ **Opting Out Costs Over Time** For $t_1, t_2 \in \mathcal{T}$ and $i \in \{1,2\}$, if $t_1 < t_2$ then $U^i((\mathbf{Opt},t_1)) > U^i((\mathbf{Opt},t_2))$.

[9] In the rest of the paper we assume that c_i is an integer. However, similar results can be obtained when c_i is any real number.

We do not make any assumption concerning the preferences of an agent for opting out versus an agreement. This enables us to consider different types of cases of opting out. Formally, there is no fixed $s \in S$ such that for every $t \in \mathcal{T}$, $U^i((s,t)) = U^i((\mathbf{Opt},t))$ as in [51].

The main factor that plays a role in reaching an agreement when agents can opt out of the negotiation is the worst agreement for agent i in a given period t which is still preferable to i than opting out in time period t. We will denote this agreement by $\hat{s}^{i,t} \in S$. If agent i will not agree to such an agreement, its opponent has no other choice but to opt out.

Definition 6 Agreements that are Preferred Over Opting Out
For every period $t \in \mathcal{T}$ and agent $i \in \mathcal{A}$ let $Possible_i^t \stackrel{\text{def}}{=} \{s^t \mid s^t \in S,\ U^i((s^t,t)) > U^i((\mathbf{Opt},t))\}$ be the set of all the possible agreements that are preferred by agent i in period t to opting out in period t. If $Possible_i^t$ is not empty, we define the agreement $\hat{s}^{i,t} \in Possible_i^t$ to be the only one that satisfies $U^i((\hat{s}^{i,t},t)) = min_{s \in Possible_i^t} U^i((s,t))$. Otherwise we define $\hat{s}^{1,t} = (-1, M+1)$ and $\hat{s}^{2,t} = (M+1,-1)$.

If Possible$_i^t$ is not empty then there will be only one minimal $\hat{s}^{i,t}$. This is because of assumption $(A1_4)$ above.

An agreement may only be reached if there is at least one agreement that both agents prefer over opting out. So, in order to reach an agreement, an agent i should prefer over opting out the worst agreement for its opponent j other than j's opting out. That is, agent i's utility from the worst agreement for agent j in a given time t that is better to j than opting out $(\hat{s}^{j,t})$ is at least equal to i's utility from the worst agreement for itself that is better than opting out $(\hat{s}^{i,t})$ (i.e., if $U^i((\hat{s}^{j,t},t)) \geq U^i((\hat{s}^{i,t},t))$). Note that by condition $(A3_4)$ if $U^i((\hat{s}^{j,t},t)) \geq U^i((\hat{s}^{i,t},t))$ then $U^j((\hat{s}^{i,t},t)) \geq U^j((\hat{s}^{j,t},t))$.

We will now introduce two additional assumptions that will ensure that an agreement will be reached.

$A5_4$ **Agreements vs. Opting Out:** For every $t \in \mathcal{T}$ $i \in \{1,2\}$, if $U^i((s,t)) > U^i((\mathbf{Opt},t))$ then $U^i((s,t-1)) > U^i((\mathbf{Opt},t-1))$.

Assumption $(A5_4)$ then indicates that if an agreement is preferred over opting out in some time period, it will also be preferred in the previous time periods over opting out. That is, the set of acceptable agreements for an agent is not increasing over time.

An additional assumption is necessary to ensure that an agreement is possible at least in the first period. That is, there is an agreement that both agents prefer over opting out.

$A6_4$ **Possible Agreement:** For all $i,j \in \mathcal{A}$, $U^i((\hat{s}^{j,0},0)) \geq U^i((\hat{s}^{i,0},0))$

$\hat{s}^{i,0}$ is the worst agreement for agent i in period 0 which is still better than opting out.

We will assume that there is some time period T in which there is no agreement that is acceptable for both agents over opting out. This time period may be viewed as a deadline.

$A7_4$ **Time Period When Agreement is not Possible:** There exists a time period T where for all $i,j \in \mathcal{A}$, $U^i((\hat{s}^{j,T},T)) < U^i((\hat{s}^{i,T},T))$. We denote the earliest of these time periods by $\hat{T}$.

4.2 Agreement is Guaranteed With No Delay

Even though the agents have the option to opt out of the negotiation in any step, if the agents use perfect equilibrium strategies, agreement will be reached without any delay. No agent will use the option of opting out since there are always agreements in the beginning of the negotiation that are better to both agents than opting out. If the agents use P.E. strategies, there is an agreement that is offered in the first time period by the first agent to make an offer, which will be preferred by its opponents over all possible future outcomes. As in the previous case (Section 3) when agents were bounded to an agreement, the main driving force for the agent to reach an agreement in this case is the cost of the negotiation time. The agents' attitudes toward opting out vs. agreements will only affect the details of the actual agreement that is reached, but won't drive any of the agents to opt out.

As the first step to proving the existence of such an agreement, we will now prove that under the above assumptions, if the negotiation has not ended in periods prior to $\hat{T}$, then an agreement will be reached in the period immediately prior to this period, i.e., in $\hat{T}-1$. The main reason is that both agents in the period prior to this period will try to avoid opting out and will agree to the worst agreement for themselves which is still better than opting out.

Lemma 2 Agreement will be Reached Prior to the Time Period When Agreement is No Longer Possible.
All the Perfect Equilibrium (P.E.) strategies of a model satisfying $A0_4 - A6_4$ satisfy the following: If it is agent 2's turn in time period $\hat{T}-1$ then using its P.E. strategy it will suggest $\hat{s}^{1,\hat{T}-1}$ and if it is agent 1's turn he will suggest $\hat{s}^{2,\hat{T}-1}$. In both cases the other party will accept the offer.

Proof: First note that by $(A6_4)$ $\hat{T} \neq 0$ and therefore $\hat{T} - 1 \in \mathcal{T}$. Now, suppose that it is agent 2's turn to make an offer at time period $\hat{T} - 1$. It is clear that agreement won't be reached after this period. Therefore, since disagreement is the worst outcome $(A0_4)$, the negotiation process will end with one of the agents opting out. Actually, since the agents prefer opting out sooner rather than later $(A4_4)$, agent 2 will opt out in the next time period. But, by $(A6_4)$ and $(A7_4)$, in time period $\hat{T} - 1$ there are still some agreements that both agents prefer over opting out (at least one). Agent 2 can choose the best agreement from its point of view and agent 1 does not have any other choice but to accept this offer. The best agreement from agent 2's point of view is $\hat{s}^{1,\hat{T}-1}$. The proof, when it is agent 1's turn to make an offer in time period $\hat{T} - 1$, is similar to this one. ∎

In the rest of the section, we assume that agent 1 is the first agent to make an offer. Since agent 1's and agent 2's positions are similar, all the results can also be proved when agent 2 is the first to make an offer.

We will now define the agreement that will be offered by an agent when it is its turn to make an offer. This agreement will be acceptable to the other agent. The intuition behind this definition is the following: in each step the agent whose turn it is to make an offer considers the possible agreement that can be reached in the following time periods. It may offer an agreement that will be better to the other agent than what that other agent can get in the next periods. However, the offer will be the worst agreement among these possible future agreements. That is, since the agents are losing over time, the first agent will offer the second agent the possible agreement in the next period minus the second agent's losses over time. The starting point is $\hat{T} - 1$ where the agreement that will be signed is clear from the previous lemma. The definition depends on whether $\hat{T}$ is even or odd.

Definition 7 Acceptable Agreements

$\hat{T}$ is even. *Suppose it is agent 2's turn to make an offer in time period $\hat{T} - 1$, (i.e, $\hat{T} - 1$ is odd). Let us define $x^{\hat{T}-1} = \hat{s}^{1,\hat{T}-1}$. For any $t \in \mathcal{T}$ $t = \hat{T} - k$, $1 < k \leq \hat{T}$, if t is even we define $x^t = (\hat{s}_1^{1,\hat{T}-1} - \frac{k}{2}c_2 + (\frac{k}{2} - 1)c_1, \hat{s}_2^{1,\hat{T}-1} + \frac{k}{2}c_2 - (\frac{k}{2} - 1)c_1)$. If t is odd we define $x^t = (\hat{s}_1^{1,\hat{T}-1} - \frac{k-1}{2}c_2 + \frac{k-1}{2}c_1, \hat{s}_2^{1,\hat{T}-1} + \frac{k-1}{2}c_2 - \frac{k-1}{2}c_1)$.*

$\hat{T}$ is odd. *In this case it is agent 1's turn to make an offer in time period $\hat{T} - 1$. Let us define $x^{\hat{T}-1} = \hat{s}^{2,\hat{T}-1}$. For any $t \in \mathcal{T}$ $t = \hat{T} - k$, $1 < k \leq \hat{T}$, if t is even we define $x^t = (\hat{s}_1^{2,\hat{T}-1} - \frac{k-1}{2}c_2 + \frac{k-1}{2}c_1, \hat{s}_2^{2,\hat{T}-1} + \frac{k-1}{2}c_2 - \frac{k-1}{2}c_1)$. If t is odd we define $x^t = (\hat{s}_1^{2,\hat{T}-1} + \frac{k}{2}c_1 - (\frac{k}{2} - 1)c_2, \hat{s}_2^{2,\hat{T}-1} + (\frac{k}{2} - 1)c_2 - \frac{k}{2}c_1)$.*

If $\hat{T}$ is even, that is, it is agent 2's turn to make an offer before the period where no agreement can be reached, there is a small advantage to agent 2. If $\hat{T}$ is odd, there is some advantage to agent 1. However, since $\hat{s}^{1,\hat{T}-1}$ and $\hat{s}^{2,\hat{T}-1}$ are quite close, the advantage in both cases is small. Also, the agent that is more patient gets a better offer.

We will show by induction on k that if the agents follow their perfect equilibrium strategies, the agent whose turn it is to make an offer will offer x^t and the other agent will accept this offer.

The main idea behind the proof is the following: Both agents prefer x^t over opting out. Furthermore, both agents prefer x^t in time period t over x^{t+1} at time period $t+1$. And x^t is the best such agreement for the agent whose turn it is to make an offer in time period t. In particular, since in $\hat{T}-1$ the agreement will be $x^{\hat{T}-1}$ (as we proved in Lemma 2), it is clear that in $\hat{T}-2$, $x^{\hat{T}-2}$ is the best option for the agent whose turn it is to make an offer in time period $\hat{T}-2$. Similarly in previous time periods.

$A8_4$ **Losses Due to Opting Out vs. Losses Resulting from Agreement** :

1. For any $t < \hat{T}$, $U^i((\hat{s}^{j,t}, t)) > U^i((\hat{s}^{i,t-1}, t-1))$.
2. For any $t < \hat{T}$, $\hat{s}_1^{1,t} - \hat{s}_1^{1,t-1} \leq \frac{1}{2}(c_2 - c_1)$ and $\hat{s}_2^{2,t} - \hat{s}_2^{2,t-1} \leq \frac{1}{2}(c_1 - c_2)$.

We note that the first part indicates that if there is at least one agreement acceptable to both sides (i.e., $U^i((\hat{s}^{j,t}, t)) > U^i((\hat{s}^{i,t}, t))$), then the agents also prefer the other agent's worst agreement that is still better to it than opting out in the next period, than opting out in the current period. If $c_1 = c_2$ then the second part (2) is always true.

In the following lemmas we will describe the proofs for the case where $\hat{T}$ is even. The proofs where $\hat{T}$ is odd are similar. We first prove that x^t is preferred by both agents over opting out.

Lemma 3 x^t **is Acceptable**
If the model satisfies assumptions $A0_4 - A8_4$ *then for any* $t \in \mathcal{T}$ $t = \hat{T} - k$, $1 < k \leq \hat{T}$, $U^i((x^t, t)) > U^i((\mathbf{Opt}, t))$.

Proof: The proof is based on backward induction on t.

Base case ($t = \hat{T} - 2$): In this case, t is even and $x^{\hat{T}-2} = (\hat{s}_1^{1,\hat{T}-1} - c_2, \hat{s}_2^{1,\hat{T}-1} + c_2)$.

We first show that the hypothesis is correct for agent 1. By $(A5_4)$, since $U^1((\hat{s}^{1,\hat{T}-1}, \hat{T} - 1)) > U^1((\mathbf{Opt}, \hat{T} - 1))$ it is also the case that $U^1((\hat{s}^{1,\hat{T}-1}, \hat{T} - 2)) > U^1((\mathbf{Opt}, \hat{T} - 2))$. But, since actions are costly $(A1_4)$, it is clear that $U^1((\hat{s}_1^{1,\hat{T}-1} - c_2, \hat{s}_2^{1,\hat{T}-1} + c_2), \hat{T} - 2) > U^1((\mathbf{Opt}, \hat{T} - 2))$.

We now show it for agent 2. By $(A8_4)$, $U^2((\hat{s}^{1,\hat{T}-1}, \hat{T} - 1)) > U^2((\hat{s}^{2,\hat{T}-2}, \hat{T} - 2)) > U^2((\mathbf{Opt}, \hat{T}-2))$. By $(A3_4)$, it is clear that $U^2((\hat{s}_1^{1,\hat{T}-1}-c_2, \hat{s}_2^{1,\hat{T}-1}+c_2), \hat{T}-2) \geq U^2((\hat{s}^{1,\hat{T}-1}, \hat{T}-1))$ and we can conclude that $U^2((\hat{s}_1^{1,\hat{T}-1} - c_2, \hat{s}_2^{1,\hat{T}-1} + c_2), \hat{T} - 2) > U^2((\mathbf{Opt}, \hat{T} - 2))$.

Induction case ($t < \hat{T} - 2$):

Suppose the hypothesis is true for any t', $t < t' \leq \hat{T} - 2$ and let $t = \hat{T} - k$.

1. If t is even $x^t = (\hat{s}_1^{1,\hat{T}-1} - \frac{k}{2}c_2 + (\frac{k}{2} - 1)c_1, \hat{s}_2^{1,\hat{T}-1} + \frac{k}{2}c_2 - (\frac{k}{2} - 1)c_1)$, which is actually $(x_1^{t+1} - c_2, x_2^{t+1} + c_2)$.

   For agent 1, by the induction hypothesis, $U^1((x^{t+1}, t + 1)) > U^1((\mathbf{Opt}, t + 1))$ and by $(A5_4)$ $U^1((x^{t+1}, t)) > U^1((\mathbf{Opt}, t))$. By $(A3_3)$ it is clear that $U^1((x_1^{t+1} - c_2, x_2^{t+1} + c_2), t) > U^1((\mathbf{Opt}, t))$.

   For agent 2, by $(A8_4)$ $U^2((\hat{s}^{1,\hat{T}-1}, \hat{T} - 1)) > U^2((\hat{s}^{2,\hat{T}-2}, \hat{T} - 2))$ and by $(A3_4)$ $\hat{s}_2^{1,\hat{T}-1} < \hat{s}_2^{2,\hat{T}-2} - c_2$ and $\hat{s}_2^{1,\hat{T}-1} + \frac{k}{2}c_2 - (\frac{k}{2}-1)c_1 < \hat{s}_2^{2,\hat{T}-2} - c_2 + \frac{k}{2}c_2 - (\frac{k}{2}-1)c_1$. It is enough to show that $\hat{s}_2^{2,\hat{T}-2} - c_2 + \frac{k}{2}c_2 - (\frac{k}{2}-1)c_1 < \hat{s}_2^{2,\hat{T}-k}$. That is, $\hat{s}_2^{2,\hat{T}-2} - \hat{s}_2^{2,\hat{T}-k} < c_2 - \frac{k}{2}c_2 + (\frac{k}{2}-1)c_1$ But, $c_2 - \frac{k}{2}c_2 + (\frac{k}{2} - 1)c_1 = (k - 2)(\frac{1}{2}(c_1 - c_2))$ and by $(A8_4)$ we can conclude that $U^2((x^t, t)) > U^2((\mathbf{Opt}, t))$.

2. If t is odd, $x^t = (\hat{s}_1^{1,\hat{T}-1} - \frac{k-1}{2}c_2 + \frac{k-1}{2}c_1, \hat{s}_2^{1,\hat{T}-1} + \frac{k-1}{2}c_2 - \frac{k-1}{2}c_1)$ which is actually $(x^{t+1} + c_1, x^{t+1} - c_1)$.

   The proof for agent 2 is similar to the proof for agent 1 when t is even. For agent 1, we need to show that $\hat{s}_1^{1,\hat{T}-1} - \frac{k-1}{2}c_2 + \frac{k-1}{2}c_1 < \hat{s}_1^{1,\hat{T}-k}$, i.e., $\hat{s}_1^{1,\hat{T}-1} - \hat{s}_1^{1,\hat{T}-k} < \frac{k-1}{2}(c_2 - c_1)$. This is clear by $(A8_4)$.

$\blacksquare$

We will now prove that x^t in period t is preferred by both agents over x^{t+1} in $t + 1$. This is due to the construction of the x^ts. For example, if t is even then agent 1 will do less in x^t than in x^{t+1}. So, it is clear that agent 1 prefers it (agent 1 also gains a period). For agent 2, it needs to do exactly what it loses over time, and therefore the utility is the same from both options.

Lemma 4 x^t **is Preferred Over** x^{t+1}
If the model satisfies assumptions $A0_4 - A8_4$ *then for any* $t \in \mathcal{T}$ $t = \hat{T} - k$, $1 < k \leq \hat{T}$, $U^i((x^t, t)) \geq U^i(x^{t+1}, t+1)$.

Proof: If t is even then $x_1^t = x_1^{t+1} - c_2$ and by $(A2_4)$ the claim is clear for agent 1. For agent 2, $x_2^t = x_2^{t+1} + c_2$ and the claim is clear by $(A3_4)$. Similarly, when t is odd. ∎

We will now state our final results for this section.

Theorem 1 Agreement will be Reached in the First Period
If the model satisfies assumptions $A0_4 - A8_4$ *and the agents follow their perfect equilibrium strategies, then*

If $\hat{T}$ *is even agent 1 will offer agent 2 in the first period* $(\hat{s}_1^{1,\hat{T}-1} - \frac{\hat{T}}{2}c_2 + (\frac{\hat{T}}{2} - 1)c_1, \hat{s}_2^{1,\hat{T}-1} + \frac{\hat{T}}{2}c_2 - (\frac{\hat{T}}{2} - 1)c_1)$ *and agent 2 will accept the offer.*

If $\hat{T}$ *is odd agent 1 will offer agent 2 in the first period* $(\hat{s}_1^{2,\hat{T}-1} - \frac{\hat{T}-1}{2}c_2 + \frac{\hat{T}-1}{2}c_1, \hat{s}_2^{2,\hat{T}-1} + \frac{\hat{T}-1}{2}c_2 - \frac{\hat{T}-1}{2}c_1)$ *and agent 2 will accept the offer.*

Proof: Clear from the above lemmas.

Example 2 *We return to the example of the newsletter deliverers. Two electronic newsletters (N1 and N2) are delivered by separate delivery services (D1 and D2). The publisher of N1 pays D1* \$200 *for the delivery of one edition of N1 to all its subscribers, and the publisher of N2 pays D2* \$225 *per delivery. Each delivery to any subscriber (i.e., a phone call to the subscriber's server) costs D1 or D2* \$1, *and each loses* \$1 *for each time period. There are M subscribers with subscriptions to both N1 and N2, and there are substantial savings to a delivery service if one or the other can deliver both newsletters. In the event that there is an agreement between D1 and D2 for joint deliveries to the M joint subscribers, then the publisher of N1 will pay D1* \$170, *and the publisher of N2 will pay D2* \$200 *(the lower prices reflect the fact that there are competing advertisers in the two newsletters, and consequently their joint delivery may detract from the sales impact of each newsletter). They must still pay* \$1 *per phone call to the server, and will lose* \$2 *over time.*
A dollar is the smallest unit of currency in this example.

Formally,
$U^1((\mathbf{Opt}, t)) = 200 - M - t$ *and* $U^1((s, t)) = 170 - s_1 - 2t$
$U^2((\mathbf{Opt}, t)) = 225 - M - t$ *and* $U^2((s, t)) = 200 - s_2 - 2t$

Suppose $M = 100$. *Then* $\hat{s}^{1,t} = (69 - t, 31 + t)$, $\hat{s}^{2,t} = (26 + t, 74 - t)$ *and* $\hat{T} = 22$.[10] *Since* $\hat{T}$ *is even, by Theorem 1 the agreement that will be reached in the first period is* $(46, 54)$.

5 One Agent Gains Over Time While the Other Loses

We now consider the case where one of the agents is gaining over time and one is losing over time. This is the usual situation when the agents are sharing a common resource, and one of them already has access to the resource and is using it during the negotiation process. That is, in each negotiation interaction, there are two agents: A which is currently using (**attached** to) the resource and W which is **waiting** to use the resource. The agents start a negotiation process on the re-division of the resource between them. A continues to use the resource until the negotiation process ends (if ever). As in the previous section, we assume that only discrete agreements can be reached. That is, $S = \{(s_A, s_W) \in I\!N^2 : s_A + s_W = M\}$[11] We assume that both agents can opt out of the negotiation.

5.1 Attributes of the Utility Functions

We will now modify the definitions of the previous sections to fit this new situation. We will need to modify some of our assumptions concerning the utility functions of the agents.

First we assume that the least preferred outcome for W is disagreement (**Disagreement**) while for A it is the most preferred outcome.

$A0_5$ **Disagreement** For each $x \in \{\{S \cup \{\mathbf{Opt}\}\} \times \mathcal{T}\}$: $U^A(x) < U^A(\mathbf{Disagreement})$ and $U^W(\mathbf{Disagreement}) < U^W(x)$. Agent A prefers disagreement over all other possible outcomes while agent W prefers any possible outcome over disagreement.

Assumption $(A0_5)$ is not a formal conclusion from assumptions $(A2_5)$ and $(A3_5)$ below, but it is well motivated by them. The reason is that we need to define the limit situation. We could define it differently as well without changing anything. For example we could assume that disagreement is the worst outcome to both agents.

Since we consider here the resource allocation problem, assumption $(A1_3)$ that asserts that the resource is valuable is also valid in this case. We denote it by $(A1_5$.)

$(A2_3)$ is no longer valid, since time is valuable only for W and not for A. Therefore we modify condition $(A2_3)$.

[10] We note that $U^i(\hat{s}^{i,t}, t) > U^i((\mathbf{Opt}, t))$ (see Definition 6).

[11] Throughout the rest of the paper, A's portion in an agreement will be written first.

$A2_5$ **Cost Over Time** For any $t_1, t_2 \in \mathcal{T}, s \in S$ and $i \in \mathcal{A}$ if $t_1 < t_2$, $U^W((s,t_1)) \geq U^W((s,t_2))$ and $U^A((s,t_1)) \leq U^A((s,t_2))$

Similarly, we modify assumption $(A3)$.

$A3_5$ **Agreement's Cost Over Time** — Each agent $i \in \{W, A\}$ has a number c_i such that: $\forall t_1, t_2 \in \mathcal{T}, s, \bar{s} \in S$, $U^i((s,t_1)) \geq U^i((\bar{s},t_2))$ iff $(s_i + c_i t_1) \geq (\bar{s}_i + c_i t_2)$, where $c_W < 0$ and $c_A > 0$.[12]

We assume that agent A gains over time ($c_A > 0$) and that agent W loses over time ($c_W < 0$), i.e., agent W prefers to obtain any given number of units sooner rather than later, while agent A prefers to obtain any given number of units later rather than sooner.

Notice that assumptions $(A1_5)$ and $(A2_5)$ are simple conclusions from assumption $(A3_5)$. We still would like to be able to distinguish between the two different properties of the utility functions. One is the desirability of the resource and the second is monotonic cost over time.

$A4_5$ **Cost of Opting Out Over Time** — for any $t \in \mathcal{T}$, $U^W((\mathbf{Opt}, t)) > U^W((\mathbf{Opt}, t+1))$ and $U^A((\mathbf{Opt}, t)) < U^A((\mathbf{Opt}, t+1))$.

W prefers opting out sooner rather than later and A always prefers opting out later rather than sooner. This is because A gains over time while W loses over time. For this reason A would *never* opt out. In the worst case A would prefer for agent W to opt out in the next period.

5.2 Agreement is Guaranteed at the Latest in the Second Period

Even though agent A prefers to continue the negotiation indefinitely, an agreement will be reached (after a finite number of periods). The reason for this is that agent W can threaten to opt out at any given time. This threat is the driving force of the negotiation process toward an agreement. If there is some agreement s that A prefers at time t over W's opting out in the next period $t+1$, then it may agree to s.

So the main factor that plays a role in reaching an agreement is the worst agreement for agent W in a given period t which is still preferable to W than opting out in time period t.

[12] We note the change in the direction of the inequality from $(A3_4)$, since here the resource is valuable, while in Section 4 we deal with costly action.

As in Section 4 we will denote this agreement by $\hat{s}^{W,t} \in S$. If agent A will not agree to such an agreement, its opponent has no other choice but to opt out.

Agent A's loss from opting out is greater than that of W. This is because A's session (of using the resource) is interrupted in the middle. Thus, we must modify assumption $A5$ to meet the current circumstances.

$A5_5$ **Range for Agreement:** For every $t \in \mathcal{T}$ $U^W((\hat{s}^{W,t}, t)) > U^W((\hat{s}^{W,t+1}, t{+}1))$, $U^W((\mathbf{Opt}, t)) > U^W((\hat{s}^{W,t+1}, t{+}1))$, and if $\hat{s}^{W,t}_A \geq 0$ then $U^A((\hat{s}^{W,t}, t)) > U^A((\mathbf{Opt}, t{+}1))$ and $U^A((\hat{s}^{W,t+1}, t+1)) > U^A((\hat{s}^{W,t}, t))$.

If there are some agreements that agent W prefers over opting out, then agent A also prefers at least one of those agreements over W's opting out even in the next period.

We assume that assumption $(A6_4)$ is still valid and denote it by $(A6_5)$. Assumptions $(A7_4)$ and $(A8_4)$ are not needed in the current situation.

We consider two cases. In the first case an agent loses less per period while waiting than it can gain per period while using the resource. In the second situation, an agent loses more while waiting for the resource than it can gain while using the resource. For this second agent, sharing a resource with others is not efficient. Therefore, it prefers to have its own private resource if possible. However, in some cases the agents don't have any choice, but to share a resource (like a road junction or another expensive resource).

We first consider the case where W *loses less over time than* A *can gain over time.* In such a case for any offer, if it is big enough, it is possible to find an offer in the future that will be better for both sides, i.e, both agents have positive total gain. Although it might appear that such an assumption will cause long delays in reaching an agreement, we will prove that in fact the delay will be at most one period since W may opt out. However, since better agreements for both parties can be found in the future, the agreement that is reached is not Pareto Optimal over time.

The intuition behind this proof is as follows. If it is not agent W's turn to make an offer in some time period t, it can always opt out and gain utility similar to that of $\hat{s}^{W,t}$ (actually, between $\hat{s}^{W,t}_W$ and $\hat{s}^{W,t}_W - 1$). So, in time period $t-1$, W will never make a better offer to A than $\hat{s}^{W,t}_A + |c_W|$, which is its benefit from opting out in the next period with the addition of W's loss over time (note that $c_W < 0$). But A will refuse such an offer, since A prefers waiting a period and offering W $\hat{s}^{W,t} \in S$. This offer will prevent W from opting out, and if W accepts the offer, A's share will be $\hat{s}^{W,t}_A + c_A$ which is better to A than $\hat{s}^{W,t}_A + |c_W|$ since $|c_W| < c_A$.

So, an agreement won't be achieved when it is W's turn to make an offer and there is still the possibility of an agreement in the next period. On the other hand if A offers W something less preferred by W than $\hat{s}^{W,t}$, W will opt out since it will never receive in any given time period in the future t' anything more than $\hat{s}^{W,t'}$, and W prefers opting out over it. So, in order to prevent W from opting out A should offer it $\hat{s}^{W,t}$ which is acceptable to W. So, if W is the first agent to make an offer (this is a reasonable assumption since A is using the resource and does not have a motive to start the negotiations), the agreement will be reached in the second period with $\hat{s}^{W,1}$. If A is the first one, agreement will be reached in the first period with $\hat{s}^{W,0}$.

The second case considers the situation where *agent W's losses over time are greater than agent A's gains.* In this model, for any agreement in period $t \in \mathcal{T}$, there is no other agreement in the future that *both* agents will prefer over this agreement. On the other hand if an agreement s in period t is small enough, one can find an agreement in a period earlier than t which both agents prefer over s in period t. According to our assumptions, this property will cause the agents to reach an agreement in the first period.

In each period in this case, if an agreement exists which agent W prefers over opting out there exists such an agreement which agent A cannot reject. The idea is the following. Agent W will accept or make an offer only if it is better for it than opting out. If A receives an offer such that there is no better agreement for it in the future, and it is also better for W than opting out, and if A prefers this offer over W opting out in the next period, it must accept this offer. Otherwise, if this agreement is rejected, W should opt out as soon as possible, since it cannot expect to do any better than opting out. But if A prefers the proposed agreement over W's opting out in the next time period, it should accept the offer.

Such an agreement, i.e., the type which will be reached in some period $t \in \mathcal{T}$ where there is still a possibility for reaching an agreement in the next time period, will be at most (from W's point of view) $\hat{s}^{W,t}$. The reason for that is that if there is still a possibility for an agreement in $t+1$, A wants to delay reaching an agreement. By offering $\hat{s}^{W,t}$ A prevents W from opting out, and gains another period of time. On the other hand, A won't accept anything worth less to it than $\hat{s}^{W,t+1}$, since A can always wait until the next period, gain a period, and reach such an agreement. Therefore, in a given time period $T+1$, A won't accept anything worth less than $\hat{s}_A^{W,T+1} + c_A + 1$. But this agreement is better to A than anything that is acceptable to W in the future. On the other hand, since W loses over time more than A can gain, this agreement is also better to W over anything it can get in the future.

The next theorem is a formal statement of the above.

Theorem 2 Agreement Will be Reached in the First or Second Period *Let $(\widehat{f}, \widehat{g})$ be a P.E. of a model satisfying $A0_5 - A6_5$. Suppose agent W is the first to make an offer.*

W loses less than A can gain. *If $|c_W| < c_A$, then $Outcome(\widehat{f}, \widehat{g}) = ((\hat{s}_A^{W,1}, \hat{s}_W^{W,1}), 1)$.*

W loses more than A can gain. *If $|c_W| > c_A$, then $Outcome(\widehat{f}, \widehat{g}) = ((\hat{s}_A^{W,1}+1+c_A, \hat{s}_W^{W,1}-1-c_A), 0)$*

If A is the first agent to make an offer then $Outcome(\widehat{f}, \widehat{g}) = ((\hat{s}_A^{W,0}, \hat{s}_W^{W,0}), 0)$.

Proof: The proof uses similar techniques to the proof of Theorem 1. The first claim is clear from the intuition above. For the second part, we will show that in any given time period T, where agreement is still possible in $T+1$, there is no agreement in the future which is better to A than $\hat{s}_A^{W,T+1} + c_A + 1$ which is also better to W than opting out in $T+1$. It is clear for time period $T+1$. Suppose $t > 1$ and suppose there is an agreement s in time $T+t$ such that $s_A + tc_A > \hat{s}_A^{W,T+1} + c_A + 1$. That is, $s_W + c_A + 1 < \hat{s}_W^{W,T+1} + tc_A$, i.e., $s_W < \hat{s}_W^{W,T+1} + c_A(t-1) - 1$. However, $|c_W| > c_A$ and therefore, $s_W < \hat{s}_W^{W,T+1} + |c_W|(t-1) - 1$ and $s_W - t|c_W| < \hat{s}_W^{W,T+1} - |c_W| - 1$ and we may conclude that $U^W((s,t)) < U^W((\hat{s}_A^{W,T+1} + 1, \hat{s}_W^{W,T+1} - 1), T+1) \leq U^W((\mathbf{Opt}, T+1))$. ∎

Example 3 *The US and Germany have embarked on a joint scientific mission to Mars involving separate mobile labs launched from a single shuttle in orbit around the planet. Each country has contracts with a number of companies for the conduct of experiments. These experiments were preprogrammed prior to launch. Arrangements were made prior to launch for the sharing of some equipment to avoid duplication and excess weight on the mission. Instructions to begin each experiment must be sent from Earth.*

The US antenna was damaged during landing, and it is expected that communications between the US and its lab on Mars will be down for repairs for one day (1440 minutes) of the planned five day duration of the mission. The US can use a weaker and less reliable backup line, but this involves diverting this line from other costly space experiments, and thus the expense of using this line is very high to the US. The US would like to share use of the German line during the one-day period so that it can conduct its planned research program. Only one research group can use the line at a time, and that line will be in use for the entire duration of the particular experiment.

A negotiation ensues between the two labs over division of use of the German line, during which time the Germans have sole access to the line, and the US cannot conduct any of its experiments (except by use of the very expensive backup). By prearrangement, the Germans are using some of the US equipment for their experiments, and are gaining \$5000 per minute. While the Germans cannot conduct any of their experiments without some US equipment, the US could conduct some of its experiments without German equipment. The US is losing \$3000 per minute during the period in which they must rely on their backup communications line. An agreement between the US and Germany to share the communications line will result in a \$1000 gain per period (minute) for each group.

If an agreement on sharing the line is not reached, the US can threaten to opt out of the arrangement. In this case, the US will be able to conduct a small portion of its experiments by using all of its equipment and no German equipment, and by using the backup communications line. The overall US gain will be \$550,000, but it will lose \$1000 per any minute of the negotiation. If the US opts out, the Germans will not be able to continue their experiments (without the US equipment) and their gain will be restricted to whatever they had gained at the point the US opted out. If the Germans opt out, they will need to pay the US \$100,000 for use of the US equipment up to that point. Note that the Germans play the role of A (attached to the communication line) and the US plays the role of W (waiting for the line). A dollar is the smallest unit of currency in this example.

Formally,

- $U^g((s,t)) = 1000s_g + 5000t$, $U^g((\mathbf{Opt}_u, t)) = 5000t$, $U^g((\mathbf{Opt}_g, t)) = 5000t - 100000$
- $U^u((s,t)) = 1000s_u - 3000t$, $U^u((\mathbf{Opt}_u, t)) = 550000 - 1000t$, $U^u((\mathbf{Opt}_g, t)) = -1000t$
- $M = 1440$

The Germans prefer any agreement over opting out.

$\hat{s}_u^{u,t} = 551 + 2t$

An agreement will be reached in the second period (period 1) with $(887, 553)$.

It should be noted that there are agreements in the future that both agents prefer over reaching the agreement $(887, 553)$ *in the second period. This is because the Germans gain more over time than the US loses over time. For example, the agreement* $(879, 561)$ *in the fourth time period (period 3) is better for both agents than* $(887, 553)$ *in the second time period. The problem is that there is no way that the US can be sure that when the fourth time period arrives, the German will offer them* $(879, 561)$. *In that time the Germans need* *to offer only* $(885, 555)$ *in order to prevent the US from opting out, and they don't have any motivation to offer more.*

6 Agents with Incomplete Information

Up to this point we have assumed that the agents have full information about each other. But in most cases, the agents do not have complete information about each other and about the environment. The incompleteness of information may be the result of different factors. For example, an agent may hide its actions from the other agents; an agent may not be able to explore the environment and may be missing information about the environment; the resources that are available to one agent are not known to the others; or one agent is not familiar with its opponent's utility function.

We will consider situations when the agents have incomplete information about each other's utility functions. The situation of incomplete information becomes even more interesting when we can expect recurring encounters between the same two agents. The agents can use information obtained in one encounter in a subsequent one. So, we will assume that there is a set of agents whose members negotiate with each other, from time to time, on sharing a resource. However, we still assume that in a given period of time no more than two agents need the same resource (assumption of a *Bilateral Negotiation* of Section 2) . When there is an overlap between the time segments in which two agents need the same resource, these agents will be involved in a negotiation process. Also the simplifying assumption 5 of the beginning of Section 2 (*No Long Term Commitments*) is still valid. That is, an agent cannot commit itself to any future activity other than the agreed-upon schedule. The relaxation of this assumption is not within the scope of this paper.

We note that a given agent may play different roles in different negotiation interactions. Sometimes, it uses the resource while another agent is also trying to use it (i.e., it plays the role of A). In other situations, an agent may need the resource while another agent is currently using it (i.e., it plays the role of W).

We assume that there is a finite set of types of agents in the environment ($Type = \{1, .., k\}$), and each has a different utility function which depends on its resource usage. The different distributions of resource usage among the agents can be due to different tasks that the agents are executing, or different configurations. For example, if all the agents are communications servers that share a common communication line then an agent that has smaller disk space will use the resource more frequently than an agent that has larger disk space. The first agent is a heavier user of the resource, while the second is the lighter user of the resource.

When $j \in Type$ plays the role of W (waiting for the resource) we denote it by W_j, and when it plays the role of A (attached to the resource) we denote it by A_j. We also assume that each agent maintains some probability belief concerning its opponent's type.[13] We denote by ϕ_j^i, where $i \in \{A, W\}$,$j \in Type$, i's probability of its opponent being of type j. We assume that $\forall i \in \{A, W\} \quad \sum_{j=1}^{k} \phi_j^i = 1$. This probability belief changes over time. We denote by $\mathcal{A}$ the set of all possible configurations of agents, i.e., $\mathcal{A} = \{W_1, W_2, ..., W_k, A_1, ..., , A_k\}$.

6.1 Sequential Equilibrium

In Sections 3, 4 and 5 we have analyzed the situation using the notion of "(subgame) perfect equilibrium." That is, we required that each agent's action be optimal for any "subgame," not just at the start of the negotiation. When there is incomplete information there is no proper subgame. In the incomplete information situation we will be talking about *sequential equilibrium* instead [29]. A sequential equilibrium includes, in addition to a profile of strategies (as in P.E.), a system of beliefs.

That is, a **Sequential Equilibrium** is a sequence of $2k$ strategies (one for each possible agent $A_1, A_2, .., A_k, W_1, ...W_k$) and a system of belief with the following properties: each agent has a belief about its opponent's type. At each negotiation step t the strategy for agent i is optimal given its current belief (at step t) and its opponents possible strategies in the S.E. At each negotiation step t each agent's belief (about its opponent's type) is consistent with the history of the negotiation. That is, the agents' belief may change over time, but only consistent with the history. We assume that each agent in a negotiation interaction has an initial probability belief about its opponent's type.

We will define these notions formally.

Definition 8 History — *for any step $t \in \mathcal{T}$ of the negotiation let $h(t)$ be the history through time step t. $h(t)$ is a sequence of t proposals and responses.*

A *strategy* for each agent as defined above specifies an action for every possible history after which it has to move. A sequence of $2k$ strategies, one for each possible agent $A_1, ...A_k, W_1, ..., W_k$, leads, from the point of view of the agents, to a probability distribution

[13] In recent work [28] we have developed a logic of probabalistic belief and time, which was integrated into the strategic model of negotiation. It was shown to be useful in cases of more complicated incomplete information.

over outcomes. For example, if agent A believes with probability ϕ_1^A that its opponent is of type 1 then A expects that with probability ϕ_1^A the outcome is determined by the strategy that is specified to A and the strategy that is specified in the sequential equilibrium to W_1. If A believes that W's type is k with probability ϕ_k^A, then it assumes that with probability ϕ_k^A the outcome will be the result of W's usage of the strategy that is specified in the sequential equilibrium for type k and its own strategy. The agents use expected utility to compare among these outcomes.

In order to state the requirement that an agent's strategy be optimal for every history, we must specify its beliefs about the other agent's type. Therefore the notion of sequential equilibrium requires us to specify two elements: the profile of strategies and the beliefs of the agents.

Definition 9 System of Beliefs — *A system of beliefs is a function $p_i(h)$ which is a probability distribution of i's opponents as a function of the history. That is, $p_i(h) = \{\phi_1^i, ..., \phi_k^i\}$ describes agent $i \in \{A, W\}$ belief about its opponent's type according to a given history of offers and counteroffers h.*

For example, suppose there are three types of agents in the environment, and suppose that before the negotiation starts A believes that with probability $\frac{1}{2}$ its opponent is of type 1, with probability $\frac{1}{4}$ it is of type 2 and with probability $\frac{1}{4}$ its opponent is of type 3. That is, $p_A(\emptyset) = (\frac{1}{2}, \frac{1}{4}, \frac{1}{4})$. Now suppose A receives an offer s from its opponent W. A may now change its beliefs. For example, it may conclude that its opponent can't be of type 3, but rather there is probability $\frac{2}{3}$ that it is of type 1 and $\frac{1}{3}$ that it is of type 2. That is, $p_A(s) = (\frac{2}{3}, \frac{1}{3}, 0)$.

We impose three conditions on the sequence of strategies and the agent's system of beliefs:

- **Sequential Rationality** — The optimality of agent i's strategy after any history h depends on the strategies of $W_1, ...W_k$ and on its beliefs $p_i(h)$. That is, agent i tries to maximize its expected utility, with regard to the strategies of its opponents and its beliefs about the probabilities of its opponent's type according to the given history. It does not take into consideration possible interactions in the future.

- **Consistency** — Agent i's belief $p_i(h)$ should be consistent with its initial belief $p_i(\emptyset)$ and with the possible strategies of its opponent. An agent must, whenever possible, use Bayes' rule to update its beliefs.

If, after any history, the strategies of the agent's opponent, regardless of its type, call for it to reject an offer and make the same counteroffer, and this counteroffer is indeed made, then the agent's beliefs remains the same as before it made the offer. If only one of the strategies of the opponent, for example, type j, specifies that the offer made by the agent may be rejected and the counteroffer s be made, and the counteroffer s is indeed made, then it believes with probability 1 that its opponent's type is indeed j.

We return to the above example, where there are three types of agents in the environment. Suppose A's original belief was $p_A(\emptyset) = (\frac{1}{2}, \frac{1}{4}, \frac{1}{4})$ as above. And suppose that the strategies of W_1, W_2 and W_3 specify that in the beginning all of them will make an offer s, then A's beliefs can't be changed if it indeed receives the offer s. However, if the strategies W_1 and W_2 specify the offer s, but W_3 specifies the offer s' then if A receives an offer s' it believes that its opponent is of type 3. That is, $p_A(s) = (0, 0, 1)$.

- **Never Dissuaded Once Convinced** — Once an agent is convinced of the type of its opponent with probability 1, or convinced that its opponent can't be of a specific type, i.e., the probability of this type is 0, it is never dissuaded from its view. The condition implies, for example, that in the above example, once agent A reaches the conclusion that its opponent is W_3 it cannot revise its belief, even if agent W subsequently deviates from W_3's strategy. From this point on A is engaged in a perfect information negotiation with agent W_3.[14]

Definition 10 Sequential Equilibrium *is a sequence of $2k$ strategies and a system of beliefs $p_i(h) i \in Type$ that satisfy the conditions of* **Sequential Rationality**, **Consistency** *and* **Never Dissuaded Once Convinced**.

6.2 Attributes of the Utility Functions

We will now change our assumptions to fit the case of incomplete information. Assumptions $A0_5 - A6_5$ are still valid. We denote them by $A0_6 - A6_6$. We only add some additional requirements to $(A4_5)$.

As in $(A4_5)$ we assume that each agent has utility with a constant cost of time c_i. We will concentrate on the cases where agent A_i $i \in Type$ gains over time ($c_{A_i} > 0$) and agent W_i loses over time ($c_{W_i} < 0$). Agent W prefers any given portion of the resource sooner rather than later, while agent A prefers any given portion of the resource later rather than

[14]See [36] for a discussion of this assumption. We leave the relaxation of this assumption for future work.

sooner. The exact values, c_{A_i} and c_{W_j} are private information. That is, agent A_i knows its private gain c_{A_i}, but may not know c_{W_j}, although it knows that it is one of k values.

We return to the situation of two agents, one A is attached to the resource, and is gaining over time, while the other W is waiting to access the resource, and is losing over time. The set of agents is $\mathcal{A} = \{W_1, W_2, ..., W_k, A_1, ..., , A_k\}$ (see the beginning of this section).

We will consider the situation where it is common knowledge that $|c_{W_k}| < |c_{W_{k-1}}| < ... < |c_{W_1}| < |c_{A_k}| < .. < |c_{A_1}|$. That is, agent W_k loses less than agent W_1 loses while waiting for the resource. Agent A_k also gains less than A_1 while using the resource. Both agents lose less while waiting than they can gain while using the resource.[15] We also assume that for any time period t, $\hat{s}_A^{W_k,t} < \hat{s}_A^{W_{k-1},t} < ... < \hat{s}_A^{W_1,t}$.[16] That is, W_k is more willing to opt out (compared with reaching an agreement) than W_1. We show that in situations that satisfy conditions $A0_6 - A6_6$, there exists a sequence of strategies that are in sequential equilibrium.[17]

6.3 Negotiation Ends in the Second Period

If the above assumptions hold and the agents use sequential equilibrium strategies, then the negotiation will end in the second period. The agents will reach an agreement in this period with high probability. The exact probability and the details of the agreement depend on agent A's initial belief. As A's belief about its opponent's type becomes more adequate, the probability that its opponent will opt out decreases.

The probability that an agreement will be reached depends also on A's type. As the difference between A's utility from agreement and its utility from opting out decreases, the probability that W will opt out decreases.

We will show that all agents that play the role of W (regardless of their type) will try to deceive their opponents, and behave as the strongest agent W_k in the first period. Agent A will ignore their offer, and will make its counteroffer in the second period, based on its initial belief and its type. In most of the cases this offer will be accepted. In the rest of the

[15] As we explained above, there are situations where an agent loses more while waiting for the resource than it can gain while using the resource. For this agent sharing a resource with others is not efficient. Therefore, it prefers to have its own private resource if possible. However, in some cases the agents have no choice but to share a resource (like a road junction or another expensive resource). Our approach is also applicable in some of these situations but common knowledge about the agents' beliefs is needed.

[16] As we defined above, for $i, j \in Type$, c_{W_i} denotes agent W_i's loss over time and c_{A_j} denotes agent A_j's gain over time. $\hat{s}^{W_i,t}$ is the worst agreement for agent W_i which is still better than opting out. The subscript A, i.e., $\hat{s}_A^{W_i,t}$ indicates A's portion of the resource in the agreement $\hat{s}^{W_i,t}$.

[17] It is enough to assume that $\hat{s}_A^{W_k,t} + |c_{W_k}| < \hat{s}_A^{W_{k-1},t} + |c_{W_{k-1}}| < ... < \hat{s}_A^{W_1,t} + |c_{W_1}|$.

paper, we will describe these results.

We first define another notion that captures the belief of an agent about how strong its opponent is.

Definition 11 The Strongest an Opponent Can Be Believed to Be
Let $p_m(h)$ be the system of beliefs of agent m after history h. Let n' be the maximal $n \in Type$ such that $\phi_n^{A_j} \neq 0$. n' is the strongest agent that m believes its opponent may be.

In the next lemma we will show the exact agreements each of the agents makes or accepts in a given time period t.

The first part of the lemma (1) indicates that W_i won't offer A anything better than its possible utility from opting out in the next period, with the addition of W_i's loss over time. It can always wait another time period and opt out. That is, W_i won't offer A anything better than its offer in the situation of full information.

In the second part of the lemma (2) we will show that A will behave toward W no better than it will toward the strongest type A believes W may be. That is, if A believes that W can't be stronger than W_j, it won't offer it more than it will offer to W_j when there is full information and W is indeed W_j.

The third part (3) of the lemma indicates that if W_i gets an offer worth less to it than opting out (less than $\hat{s}^{W_i,t}$), it will really opt out. We will show that if W_i rejects this offer it won't receive any future offers better than this one. But it prefers opting out over reaching this agreement. On the other hand, if W_i does receive an agreement at least as good as it can get from opting out, it should accept it. W_i won't be offered any agreement better than that (part (4) of the lemma).

Lemma 5 Agreements that are Accepted and Agreements that are Rejected by the Agents
Suppose agent W is of type $i \in Type$ and A is of type $j \in Type$, and the agents' utility functions satisfy $A1_6 - A6_6$, $\hat{s}_A^{W,t} - \hat{s}_A^{W,t+1} < c_{A_k}$. Let $n_A \in Type$ be the strongest type A believes its opponent can be (as defined in Definition 11). If both W and A use their sequential equilibrium strategies then the following holds:

1. **The best offer to A that may be made by W_i:**
   W_i will not offer A in step t more than $\hat{s}_A^{W_i,t+1} + |c_{W_i}|$.

2. **The best agreement to agent W that may be made or accepted by A_j:** *A_j will not accept anything less than $\hat{s}_A^{W_{n_A},t+1} + c_{A_j}$ in step t and won't offer anything more than $\hat{s}^{W_{n_A},t}$.*

3. **When W_i will opt out** *If in step t the offer that A_j makes to W_i is less than $\hat{s}_W^{W_i,t}$, then W_i opts out in step t.*

4. **The offers that will be accepted by W_i:** *If in step t, A_j makes an offer s such that $s_W \geq \hat{s}_W^{W_i,t}$, W_i accepts the offer.*

Proof:

1. It is clear that any type of W won't offer A more than $\hat{s}_A^{W_i,t} + |c_{W_i}|$ since it can always wait until the next time period, opt out, and achieve a better outcome.

2. When it is A's turn to make an offer in a given time period t, and it believes that n_A is the strongest type W may be, then it will never offer anything better for W than $\hat{s}^{W_{n_A},t}$. This offer will prevent any type of W that A believes W can be, from opting out. This is the main goal of A; if W rejects its offer but doesn't opt out, A earns another time period of using the resource.

   If A is offered anything less than $\hat{s}_A^{W_{n_A},t+1} + |c_{A_j}|$ amount of the resource in time period t it should reject this offer; it can always wait another time period, offer W $\hat{s}^{W_{n_A},t+1}$ and since it gains over time c_{A_j}, its utility will be at least as large as the utility of $\hat{s}_A^{W_{n_A},t+1} + |c_{A_j}|$ at time t.

3. This will be proved by induction on the number of types ($|Type|$). We note that $U^{W_i}((\mathbf{Opt},t)) > U^{W_i}((\hat{s}^{W_i,t+1}, t+1))$. This is because we deal with a discrete case and by assumption ($A5_6$).

   Base case: (only two types) k=2

   If there are only two types, it is clear by the second part of the lemma (2) that in any future time period t' A won't offer W more than $\hat{s}^{W_2,t'}$. However W_2 prefers opting out now over the possibility of getting $\hat{s}^{W_2,t'}$ in future time periods t'. The only way for A to prevent W_2 from opting out now is by offering at least $\hat{s}^{W_2,t}$, which is the worst agreement to W_2 that is still better than opting out. That is, if W_2 is offered anything less than $\hat{s}^{W_2,t}$ at time t it will opt out.

Suppose A offers less than $\hat{s}^{W_1,t}$. It is clear that in such a situation W_2 will opt out since $\hat{s}_W^{W_1,t} < \hat{s}_W^{W_2,t}$. So, if W doesn't opt out after receiving such an offer, it will be clear to A that its opponent is W_1. So, by (2) A won't offer to W in the future (t') anything better to W than $\hat{s}^{W_1,t'}$, but W_1 prefers opting out now over $\hat{s}^{W_1,t'}$ in time t'.

k=k'+1 Suppose the assumption is correct for $|Type| = k'$, and a new type i is added. If A offers something less than $\hat{s}^{W_i,t}$, all the other types of W that are stronger than i opt out, by the induction hypothesis. Therefore, if W_i won't opt out, A will know that its opponent is at most i, and won't offer anything better in the future (t') than $\hat{s}^{W_i,t'}$ (by (2)).

4. Similar to (3)

▮

Based on this lemma we prove that all agents that play the role of W, regardless of their type, will behave as the strongest type when it is their turn to make an offer. From (1) of lemma 5, it is clear that agent W_i won't offer anything more to A than $\hat{s}_A^{W_i,t+1} + |c_{W_i}|$. If so, if there is an agent that offers more, A can conclude that its type is weaker than i. But in such a case, A is better off waiting and offering W an agreement on a scale that it might make if its opponent was weak. So, it isn't worth it for an agent in the first period to reveal that it is weak. It can only lose from that. Therefore, all the types of agent W behave as the strongest one.

When it is A's turn to make an offer, it should try to maximize its expected utility. So, it needs to calculate for which agreement, according to its beliefs, its expected utility will be the highest. This is, by taking into account that if it offers its opponent an agreement that fits W_i ($\hat{s}^{W_i,t}$), and its opponent is stronger than i, it will opt out. If it is of type i or weaker than i it will accept the offer.

Lemma 6 W will Pretend to be Strong and A will Behave According to its Expected Utility

Suppose, agent W is of type $i \in Type$ and A is of type $j \in Type$, and the agents' utility functions satisfy $A0_6 - A6_6$, $\hat{s}_A^{W,t} - \hat{s}_A^{W,t+1} < c_{A_k}$. If both W and A use their sequential equilibrium strategies then the following properties hold:
(i) An agent of type $i \in Type$ will accept any offer greater or equal to $\hat{s}_W^{W_i,t}$.
(ii) All agents of type W, regardless of their types, will offer in any time period t $\hat{s}_A^{W_k,t+1} + |c_{W_k}|$ amount of the resource.

(iii) Suppose A has a probability belief of $(\phi_1^A, ..., \phi_k^A)$, where $\phi_1^A + ... + \phi_k^A = 1$. Let $Expect(\hat{s}^{W_i,t}) = (\phi_1 + ... + \phi_i)U^A((\hat{s}^{W_i,t}, t)) + (\phi_{i+1} + ... + \phi_k)U^A((\mathbf{Opt}, t))$. Let $\hat{i} \in Type$ such that $Expect(\hat{s}^{W_{\hat{i}},t})$ is maximal over any $i \in Type$. A will offer $\hat{s}^{W_{\hat{i}},t}$ in time period t when it is its turn to make an offer.

Proof:
(i) It is clear by (4) of Lemma 5.
(ii) (by induction on the number of types)
Base Case ($k = |Type| = 2$).
By (1) of Lemma 5 it is clear that W_2 will not offer anything more to A than $s_A^{W_2,1} + |c_{W_2}|$ which will be rejected by A according to (2) of Lemma 5. So if W_1 will offer something better than $s_A^{W_2,1} + |c_{W_2}|$, A can conclude that its type is 1, i.e., $n_A = 1$. But, by (2) of the lemma, it will reject the offer, and in the next period will offer W_1 nothing more than $\hat{s}^{W_1,1}$. Therefore, W_1 should prefer to pretend to be W_2.
Inductive case ($k = |Type| = k' + 1$).
Suppose the induction hypothesis is correct for k' number of types. And suppose another type i' is added which is weaker than the previous types. From the induction hypothesis it is clear that all of them, if they play the role of W, will pretend to be strong, therefore if $W_{i'}$ will behave differently, A will know who it is, reject its offer, and won't offer it anything better than $\hat{s}^{W_{i'},t+1}$ in the future. So, if i plays the role of W, it should pretend to be strong.
(iii) By (3) of Lemma 5, it is clear that if A will offer $s^{W_i,1}$ in period 1, all the agents that are stronger than i will opt out, while the others will accept it. So, A calculates its expected utility from all its options, and chooses the best option for itself. ∎

We will now state our final results for this section.

Theorem 3 Either an Agreement Will Be Reached in the Second Period, or W Will Opt Out

Suppose agent W is of type $i \in Type$ and A is of type $j \in Type$, and the agents' utility functions satisfy $A0_6 - A5_6$, $\hat{s}_A^{W,t} - \hat{s}_A^{W,t+1} < c_{A_k}$. Let $\hat{i} \in Type$ such that $Expect(\hat{s}^{W_{\hat{i}},2})$ is maximal over any $i \in Type$ (where Expect is defined as in Lemma 6 (iii)).

If both W and A use their sequential equilibrium strategies then in the first period (period 0) all types of agents playing the role of W will offer A $\hat{s}_A^{W_k,1} + |c_{W_k}|$ amount of the resource. A will reject the offer. In period 1 A will offer $\hat{s}^{W_{\hat{i}},1}$. If W is at least of type $\hat{i}$ it will accept the offer. Otherwise, it will opt out.

Proof: Clear by Lemma 6. ∎

We would like to indicate that A_j may revise its beliefs about its opponent's type after each negotiation session. If in the second period, W accepts an offer equal to $\hat{s}_W^{W_i,1}$, A concludes that W is at most of type i. If A offers $\hat{s}_W^{W_i,1}$ and W opts out of the negotiation A concludes that it is of type greater than i. Using this additional information about W, agent A can update its beliefs about other agents. For example, if A knows that there is at most one agent in the system that is of type 1, and it finds out that its opponent from the previous interaction is of type 1, then it can adjust its beliefs about the types of other agents. The updated belief will be used in future interactions.

This is the only case, among the ones that we have studied, in which one of the agents may opt out. However, as more interactions occur, more information about one another is collected, and less opting out will occur in the future.

Example 4 *We return to the example of the mission to Mars. Suppose that each of the labs (agents) on Mars does not know the exact details of the contracts the other has with companies. There are two possibilities for the contracts: high (h) and low (l). If the type of contracts the German's hold is h, then their utility functions are similar to those of Example 3. They gain \$5000 per minute during the negotiation and gain \$1000 per minute when they share the line with the US. If the US also holds contracts of type h, then their utility functions are also similar to those of Example 3. The US loses \$3000 per minute during the negotiation period and gains \$1000 per minute when sharing the line with the Germans. If the US opts out the overall US gain will be \$550,000, but they will also lose \$1000 per minute during the negotiation.*

However, if the German contracts are of type l, then they only gain \$4000 per minute while using the line by themselves. The US losses while negotiating if their contracts are type l are only \$2000 per minute. But if the US opts out, their overall gain is only \$450,000. They still negotiate for the usage of the German line in the next 24 *hours (i.e.,* $M = 1440$*) from the time the negotiation ends.*

Formally, let $s \in S$, $t \in \mathcal{T}$,

Type h	*Type l*
$U^{g_h}((s,t)) = 1000s_g + 5000t$	$U^{g_l}((s,t)) = 1000s_g + 4000t$
$U^{g_h}((\mathbf{Opt}_u,t)) = 5000t$	$U^{g_l}((\mathbf{Opt}_u,t)) = 4000t$
$U^{g_h}((\mathbf{Opt}_g,t)) = 5000t - 1000$	$U^{g_l}((\mathbf{Opt}_g,t)) = 40t - 1000$
$U^{u_h}((s,t)) = 1000s_u - 3000t$	$U^{u_l}((s,t)) = 1000s_u - 2000t$
$U^{u_h}((\mathbf{Opt}_u,t)) = 550000 - 1000t$	$U^{u_l}((\mathbf{Opt}_u,t)) = 450000 - 1000t$
$U^{u_h}((\mathbf{Opt}_g,t)) = -1000t$	$U^{u_l}((\mathbf{Opt}_g,t)) = -1000t$
$\hat{s}^{u_h,t} = (889 - 2t, 551 + 2t)$	$\hat{s}^{u_l,t} = (989 - t, 451 + t)$

Let us assume that Germany (playing the role of A) is of type h and the US (playing the role of W) is of type l. We denote them by g_h and u_l. We will consider two cases. Suppose g_h believes that with probability 0.5 *its opponent is of type h and with probability* 0.5 *its opponent is of type l, i.e., $\phi^g_l = 0.5$ and $\phi^g_h = 0.5$. According to Theorem 3, in the first period, u_l will pretend to be of type h and will offer g_h $(\hat{s}^{u_h,1}_g + |c_{u_h}|, \hat{s}^{u_h,1}_u - |c_{u_h}|) = (890, 550)$. g_h will reject the offer. In the second time period, g_h compares between offering W* $(887, 553)$*, which will be accepted by both types, and offering* $(988, 452)$ *which will be accepted by type l, but after such an offer, if W is of type h, it will opt out. Since its expected utility from offering* $(887, 553)$ *is higher, it makes this offer which is accepted by u_l.*

However, suppose g_h believes only with probability 0.1 *its opponent is of type h, and with probability* 0.9 *its opponent is of type l. The behavior of u_l in the first period is similar to the previous case. It pretends to be h. However, in the second period, g_h's expected utility from* $(988, 452)$*, is higher than* $(887, 553)$ *and therefore it makes this offer to W, which is accepted by u_l.*

7 Negotiation Mechanism for Multiple Fully Informed Agents

Up to this point we have assumed that only two agents participate in each interaction (Assumption 1, Bilateral Negotiation, in Section 2). We will now relax this assumption by extending the framework of Section 4. We note that we assume in this section that the agents have full information.

We assume that a set of agents wants to satisfy a goal. All agents can take part in satisfying the goal, but they all need to agree on the schedule. There are no side payments, i.e., no private deals can be reached among the agents. An additional option which we do not deal with in this paper involves one of the agents opting out, and the remaining agents reaching an agreement. In the rest of the section when we mention *several* agents we mean more than two.

Example 5 *We return to the example of the newsletter deliverers. Here there are several electronic newsletters (more than two), that are delivered by separate delivery service agents. The delivery is done by phone (either by fax machines or electronic mail). The expenses of the agents depend only on the number of phone calls. There are several subscribers that subscribe to all the newsletters. All the delivery agents negotiate over the distribution of the common subscriptions. Each of the agents can opt out of the negotiations and deliver all of its own newsletters by itself. The agents are paid according to the time of the delivery (the faster the better).*

As in previous sections, we denote the set of agents by $\mathcal{A}$. We denote the number of agents (i.e., $|\mathcal{A}|$) by n and we attach to each agent an integer between 1 and n.

Definition 12 Agreement:
An agreement is a tuple $(s_1, .., s_n)$, where $s_i \in I\!N$ and $s_1 + ... + s_n = M$. s_i is agent i's portion of the work.[18]

We will now modify the negotiation procedure to fit the multi-agent interaction. As in previous sections, the agents can take actions only at certain times in $\mathcal{T}$. In each period $t \in \mathcal{T}$ one of the agents, say i, proposes an agreement to all the other agents. Each of the agents either accepts the offer (chooses **Yes**) or rejects it (chooses **No**), or opts out of the negotiation (chooses **Opt**). If the offer is accepted by all the agents, then the negotiation ends and the agreement is implemented. Also, opting out by one of the agents ends the negotiations. After a rejection by at least one agent, another agent must make a counter offer, and so on. That is, this mechanism provides each of the agents with a veto power on the agreements that will be reached. However, we will show that even under this extreme assumption, agreement is guaranteed in the first time period.

We require that the agents always make offers in the same order, i.e., if $|\mathcal{A}| = n$ the first agent makes offers in time periods 0, n, $2n$..., the second agent makes offers in periods 1, $n + 1$..., etc. We do not make any assumption about who begins the negotiations, i.e., who makes the first offer and who is the second agent to make an offer. So, without loss of generality we assume that the agents are numbered according to the order in which they make the offers, i.e., the first agent is denoted by 1, the second by 2, etc.

[18] A similar definition can be given concerning a division of resources.

Definition 13 Negotiation Strategies:
A strategy is a sequence of functions. The domain of the ith element of a strategy is a sequence of agreements of length i and its range is the set $\{\textbf{Yes},\textbf{No},\textbf{Opt}\} \cup S$. *We first define a strategy f for an agent i which is the first agent to make an offer.*

Let $f = \{f^t\}_{t=0}^{\infty}$, *where* $f^0 \in S\}$, *for* $t = kn, k \in \mathcal{T}$ $f^t : S^t \to S$, *and for other* $t \in \mathcal{T}$, $f^t : S^t \times S \to \{\textbf{Yes},\textbf{No},\textbf{Opt}.\}$ *We denote by F the set of all strategies of the agent who starts the bargaining.*

Similarly, the strategies for the other agents are defined.

7.1 Attributes of the Utility Functions

We will now modify the assumptions of the previous sections to fit the multiagent situation. We note that we are dealing with the task distribution case.

$A1_7$ **Actions Are Costly:** For all $t \in \mathcal{T}, r, s \in S$ and $i, j \in \mathcal{A}gent, i \neq j$: $r_i > s_i \Rightarrow U^i((r,t)) < U^i((s,t))$. For agreements that are reached within the same time period each agent prefers to perform a smaller portion of the labor.

However, the utilities of the agents from agreements in which their parts are equivalent may be different. That is, $r_i = s_i \not\rightarrow U^i((r_i,t)) = U^i((s_i,t))$. Other parameters, such as the quality of the performance of the other agents, may also play a role.

We first consider the case that all agents lose over time. Assumption A2 of sections 3 and 4 is still valid. That is:

$A2_7$ **Time is valuable** For any $t_1, t_2 \in \mathcal{T}, s \in S$ and $i \in \mathcal{A}$ if $t_1 < t_2$, $U^i((s,t_1)) \geq U^i((s,t_2))$ and $U^i((\textbf{Opt},t_1)) \geq U^i((\textbf{Opt},t_2))$.

Assumption $(A3)$ of Section 3 and Section 4 is not valid.

We also need to modify the definition of $\hat{s}^{i,t}$ in order to be able to compare between agreements and opting out.

Definition 14 Agreements that are Preferred Over Opting Out:
For every $t \in \mathcal{T}$ *and* $i \in \mathcal{A}$ *we define* $\bar{S^i}^t = \{s | s \in S \text{ s.t. } U^i((s,t)) > U^i((\textbf{Opt},t))\}$. *We denote by* $\bar{\mathcal{S}}^t$ *the set of agreements that are preferred by* all *agents over opting out, i.e.,* $\bar{\mathcal{S}}^t = \cap_{i\in\mathcal{A}} \bar{S^i}^t$

Suppose all agents are losing over time. We assume that all agents prefer to reach a given agreement sooner rather than later, i.e., assumption A2 is still valid. We also assume that all agents prefer to opt out sooner rather than later, that is assumption A4 is valid. We also assume that if an agent prefers an agreement over opting out in some period t, then it prefers the same agreement in time period t' prior to t over opting out in t'. Formally:

$A4_7$ **Opting Out Costs Over Time** For $t_1, t_2 \in \mathcal{T}$ and $i \in \mathcal{A}$, if $t_1 < t_2$ then $U^i((\mathbf{Opt}, t_1)) > U^i((\mathbf{Opt}, t_2))$. If $U^i((s,t)) > U^i((\mathbf{Opt}, t))$, then for any $t' \in \mathcal{T}$ such that $t' < t$, $U^i((s,t')) > U^i((\mathbf{Opt}, t'))$. If $U^i((s,t)) > U^i((\mathbf{Opt}, t+1))$, then for any $t' \in \mathcal{T}$ such that $t' < t$, $U^i((s,t')) > U^i((\mathbf{Opt}, t'+1))$.

It is clear that under this assumption, for any $t \in \mathcal{T}$ $\cap_{i \in \mathcal{A}} \bar{S^i}^{t+1} \subseteq \cap_{i \in \mathcal{A}} \bar{S^i}^t$.

We will concentrate on the cases where there is an agreement in the first time period which is acceptable to all agents and that there is a time period where there is no agreement that all agents prefer over opting out. This is stated in the next assumption.

$A5_7$ **Possible Agreements** $\cap_{i \in \mathcal{A}} \bar{S^i}^0 \neq \emptyset$. There exists $T \in \mathcal{T}$ such that $\bar{S}^T = \emptyset$. We denote the minimal time period among these time periods by $\hat{T}$.

7.2 Agreement is Guaranteed With No Delay

We are able to show that the results of Section 4 are also valid when there are more than two agents in the environment and when all agents have veto power. As in the bilateral case, if the agents use perfect-equilibrium strategies, agreement will be reached without any delay. The main driving force of the agent reaching agreement in this case is the cost of the negotiation time. The agents' attitudes toward opting out v.s. agreements will only affect the details of the actual agreement that will be reached, but won't drive any of the agents to opt out.

We will first show that in such a case if the game has not ended in prior periods, then an agreement will be reached in the period prior to that in which there is no agreement acceptable to all agents, i.e., in period $\hat{T}-1$.

Lemma 7 Agreement will be Reached Prior to the Time Period When Agreement is No Longer Possible (N agent version)
Suppose the model satisfies $A0_7 - A2_7, A4_7 - A5_7$. If the agents are using their perfect equilibrium strategies, and the negotiation process is not over until time $\hat{T}-1$ and it is agent i's

turn to make an offer, then it will offer $\hat{s} = max_{U^i}\bar{S}^{\hat{T}-1}$ and all the other agents will accept the offer.[19]

Proof: In period $\hat{T}$ and later, there is no agreement that is acceptable to all the agents. Therefore, the only possible outcome after that time period is either opting out or disagreement. Since disagreement is the worst outcome ($A0_7$) and the agents prefer opting out sooner rather than later, at least one of the agents will opt out at time period $\hat{T}$. But all the agents prefer an agreement from $\bar{S}^{T-1}$ over opting out in the next period. Since it is i's turn, it can choose the best agreement from its point of view, offer it, and all the agents will accept it. ▮

Now, we will show that in each time period less than $\hat{T}$ there is always a set of possible agreements, acceptable to all the agents. The agent whose turn it is to make an offer, should choose the best of these agreements according to its utility function and make this offer.

We first define the sets of acceptable agreements by induction of t. In the period before $\hat{T}$, this set contains only $\hat{s}$. In the prior period ($\hat{T} - 2$) it includes all the agreements that are preferred in this time period by all agents over opting out and over $\hat{s}$ in $\hat{T} - 1$. From this set the best agreement for the agent whose turn it is to make an offer is chosen. In computing the acceptable agreements set in the prior period $\hat{T} - 3$ this value is used as the basis, similar to $\hat{s}$ before, i.e. the acceptable agreements in $\hat{T} - 3$ are those that are better to the agents than this value and also better than opting out. Similarly for prior periods.

Definition 15 Acceptable Agreements *Let $x^{\hat{T}-1} = \hat{s}$ (where $\hat{s}$ is as defined in Lemma 7). For each $t \in \mathcal{T}, t < \hat{T}-1$ let X^t include all the agreements that satisfy the following condition: $s \in X^t$ iff $s \in \bar{S}^t$ and for any $j \in \mathcal{A}$ $U^j((s,t)) \geq U^j((x^{t+1}, t+1))$. If it is i's turn to make an offer in time period t, we define $x^t = max_{U^i}X^t$.*

This definition is sound, since X^t is not empty for any time period before $\hat{T} - 1$. We will prove this in the next lemma. The intuition behind the proof is that x^{t+1} always belongs to X^t, since the agents lose over time, and if an agreement is preferred over opting out, at a given time period, it is also preferred over opting out in previous time periods.

[19]For simplification, we will assume that only one such maximal agreement exists. This will be the case either if the only factor determining the utility for an agent is its own portion of the task (see example 6), or if the quality of the performance of the other agents yields a different utility for each agreement. Without this uniqueness assumption, if there are several agreements which have the same maximal utility for agent i, it is difficult for the other agents to predict which offer agent i will make. These agreements may have different utilities to the other agents. We can assume that in such situations, agent i chooses any one of the maximal agreements with equal probability, and that the other agents behave according to their expected utilities.

Lemma 8 Acceptable agreements do exists: *If the model satisfies conditions $A0_7 - A2_7, A4_7 - A5_7$ then for $t \in \mathcal{T}, t < \hat{T} - 1$, $X^t \neq \emptyset$.*

Proof: We will show by backward induction on t that for all $t < \hat{T}$ $X^t \neq \emptyset$.
Base case ($t = \hat{T} - 1$):
By ($A2_7$) $\forall i \in \mathcal{A}, U^i((\hat{s}, \hat{T}-2)) > U^i((\hat{s}, \hat{T}-1))$ and by ($A4_7$) it is clear that $U^i((\hat{s}, \hat{T}-2)) > U^i((\mathbf{Opt}, \hat{T}-2))$, and therefore $\hat{s} \in A^{\hat{T}-2}$.
Inductive Case ($t < \hat{T} - 1$):
By the induction hypothesis, $X^{t+1} \neq \emptyset$. Therefore, x^{t+1} is well defined. But, by ($A2_7$) and ($A4_7$,) similarly to the base case, it is easy to see that $x^{t+1} \in X^t$ ▮

We will show now that in any time period all the agents will accept x^t, and the agent whose turn it is to make an offer will also offer x^t. The intuition behind this is that x^t is preferred by the agents over opting out, and it is better than any agreement that can be reached in the future.

Lemma 9 x^t **is offered and accepted** *If the model satisfies conditions $A0_7 - A2_7, A4_7 - A5_7$ then in any time period $t < \hat{T} - 1$, the agents will accept any offer $s \in X^t$ and the agent whose turn it is to make an offer will offer x^t.*

Proof: The proof is by backward induction of t.
Base case ($t = \hat{T} - 2$):
By Lemma 7, it is clear that the agents won't reach any agreement in the future better than $x^{\hat{T}-1}$. However, by the definition of X^t (Definition 15), any agreement in $X^{\hat{T}-2}$ is better for the agents than $x^{\hat{T}-1}$ in time period $\hat{T} - 1$, and than opting out at $\hat{T} - 2$. Therefore, they should accept these offers.

On the other hand, any agreement that does not belong to $X^{\hat{T}-2}$ won't be accepted by at least one of the agents, since it will prefer to wait another period and receive $\hat{s}$ or even opt out. But since i, similar to the other agents, prefers the agreements of $X^{\hat{T}-2}$ over this possibility, he should offer an agreement from this set. However, since it is its turn to make an offer, it has the opportunity to choose the best one from its point of view.
Inductive Case ($t < \hat{T} - 2$):
By the induction hypothesis, if an agreement isn't reached in this time period, the outcome of the negotiation process will be $(x^{t+1}, t + 1)$. But, the agreements of X^t at time period t are preferred by all agents over $(x^{t+1}, t + 1)$ and over opting out at t; the proof proceeds as in the base case. ▮

We summarize our results by the following theorem.

Theorem 4 *If the model satisfies conditions $A0_7 - A2_7, A4_7 - A5_7$, and the agents use their perfect equilibrium strategies, then in the first time period agent* 1 *will offer x^0, and all other agents will accept the offer.*

Proof: Clear, by Lemma 9. ∎

Example 6 *We return to the example of the newsletter deliverers. Three electronic newsletters (N1, N2 and N3) are delivered by separate delivery services (D1, D2 and D3). The payment arrangements for D1 and D2 are as previously discussed in example 2, i.e., the publisher of N1 pays D1* \$200 *per delivery of one edition, and the publisher of N2 pays D2* \$225 *per delivery of one edition. The publisher of N3 pays D3* \$250 *per delivery of one edition. As was the case for D1 and D2, each delivery to a given subscriber (i.e., a phone call to this subscriber's server) also costs D3* \$1, *and each loses* \$1 *for each time period. There are M subscribers with subscriptions to all newsletters (i.e., N1, N2 and N3), and as in example 2 there are substantial savings to a delivery service if one of the agents can deliver all newsletters to the same subscribers. If an agreement among D1, D2 and D3 for joint deliveries to the M joint subscribers is reached, then the publisher of N3 will pay D3 only* \$215 *per delivery of an edition, and as in the previous example, in such an event the publisher of N1 will pay D1* \$170, *and the publisher of N2 will pay D2* \$200. *They must still pay* \$1 *per phone call to the server, and will lose* \$2 *for any negotiation time period. Notice that in this example, only the number of phone calls to the subscribers made by a delivery agent plays a role in its payments and not the distribution of the rest of the subscribers between the other two agents.*

Formally,
$U^1((\mathbf{Opt}, t)) = 200 - M - t$ *and* $U^1((s, t)) = 170 - s_1 - 2t$
$U^2((\mathbf{Opt}, t)) = 225 - M - t$ *and* $U^2((s, t)) = 200 - s_2 - 2t$
$U^3((\mathbf{Opt}, t)) = 250 - M - t$ *and* $U^3((s, t)) = 215 - s_3 - 2t$
Suppose $M = 100$. Then $\hat{s}_1^{1,t} = 69 - t$, $\hat{s}_2^{2,t} = 74 - t$ and $\hat{s}_3^{3,t} = 64 - t$. Note that for all $i \in \{1, 2, 3\}$, $\hat{s}^{i,t}$ is not unique in this case.

$\hat{T} = 36$ and it is D3's turn to make an offer in the time period prior to $\hat{T}$. In this period, D1 is willing to deliver up to 34 *newsletters, if an agreement will be reached (i.e., $\hat{s}_1^{1,35} = 34$) and D2 is willing to deliver up to 39 newsletters (i.e., $\hat{s}_2^{2,35} = 39$). So, $x^{35} = (34, 39, 27)$. It is easy to compute, that whenever it is D1's turn to make an offer (t is divided by 3), $x^t = (31, 38, 31)$, when it is D2's turn to make an offer, $x^t = (35, 36, 29)$ and when it is D3's*

Type & Section	Opting Out	Num Agents	Full Info	Who loses	$\hat{T}$	c_1 ? c_2	Results
Resource 3	No	2	Yes	Both	–	$c_1 > c_2$	$((1,0),0)$
Resource 3	No	2	Yes	Both	–	$c_1 < c_2$	$((c_2, 1-c_2), 0)$
Task 4	Yes	2	Yes	Both	even	—	$((\hat{s}_1^{1,\hat{T}-1} - \frac{\hat{T}}{2}c_2 + (\frac{\hat{T}}{2}-1)c_1,$ $\hat{s}_2^{1,\hat{T}-1} + \frac{\hat{T}}{2}c_2 - (\frac{\hat{T}}{2}-1)c_1), 0)$
Task 4	Yes	2	Yes	Both	odd	—	$((\hat{s}_1^{2,\hat{T}-1} - \frac{\hat{T}-1}{2}c_2 + \frac{\hat{T}-1}{2}c_1,$ $\hat{s}_2^{2,\hat{T}-1} + \frac{\hat{T}-1}{2}c_2 - \frac{\hat{T}-1}{2}c_1), 0)$
Resource 5	Yes	2	Yes	1 (W)	—	$\lvert c_W\rvert \leq c_A$	$(\hat{s}^{W,1}, 1)$
Resource 5	Yes	2	Yes	1 (W)	—	$\lvert c_W\rvert \geq c_A$	$((\hat{s}_A^{W,1} + 1 + c_A, \hat{s}_W^{W,1} - 1 - c_A), 0)$
Task 7	Yes	$N > 2$	Yes	all	—	—	$(x^0, 0)$
Resource 6	Yes	2	No	1 (W)	—	$\lvert c_W\rvert \leq c_A$	$(\hat{s}^{W_i,1}, 1)$ where $Expect(\hat{s}^{W_i,2})$ is maximal OR W opts out in period 1

Figure 2: Summary of Results. $\hat{T}$ of column 6 is the earliest time period, in which there is no agreement preferred by both agents over opting out. c_i of column 7 is the constant cost of delay. $\hat{s}^{i,t}$ is the worst agreement for agent i in a given period t which is still preferable to i than opting out in time period t. x^0 of row 7 in column 8 is defined in Definition 15. $Expect$ is defined in Lemma 6

turn to make an offer (prior to time period 35*),* $x^t = (33, 40, 27)$. *Therefore, in the first time period (0), D1 will offer (31,38,31), and its opponents will accept its offer.*

8 Evaluation of the Results

We will analyze the results of the paper using the criteria of Section 1.3. A summary of the results appears in Figure 2.

- **Distributed** — In all the situations we analyzed there is no central unit that is involved in the inter-agent encounters. The agents negotiate to reach an agreement.
- **Instantaneously** — Conflict are resolved without delay. In most of the situations we dealt with, an agreement will be reached in the first time period. This includes the situations of Sections 3 and 4 where two agents negotiate and both are losing over time, and the situations of Section 7 where there are multiple agents that negotiate and all of them are losing over time. However, also when one of the agents is gaining over time (Section 5), but the loser W loses more than the gainer (A) gains over time, agreement

will be reached in the first period. In the second case of Section 5, when W loses less than A gains (which is the more common situation), agreement will be reached in the second period. Similarly, even if there is incomplete information (Section 6) the negotiation will end at the second period.

- **Efficiency:** In almost all the cases that we analyzed, the agents reach an agreement and conflict is avoided. Furthermore, in the cases of task distribution (Sections 4 and 6) and in the cases where the resource is not in use during the negotiation process (Sections 3), agreement will be reached in the first period. So, in the case of two agents that are waiting for the resource, since agreement will be reached in the first period, there won't be a deadlock.

  The only situation where the negotiation process may end up with an agent opting out, is in the case of incomplete information (Section 6). Also in this situation, in most of the cases the negotiation process will also end with an agreement, but in some cases it will end with W opting out. This depends on the accuracy of A's beliefs about its opponent and A's utility from opting out v.s. reaching an agreement. However, after each negotiation interaction, A may learn more about its opponent's type and update its belief. The beliefs of the agents about the types of other agents become more and more accurate with each interaction. Therefore the probability that the resource will be damaged is decreasing over time and the **efficiency** of the system is increasing.

  In most of the cases we dealt with, the agreement that is reached is Pareto optimal. That is, there is no other agreement that all the agents prefer over the one that they reached. However, in the case of resource allocation (Sections 5 and 6) where the agent that loses over time loses less than the agent that gains over time, the agreement that is reached is not always Pareto optimal. If A's share in the agreement is big enough, there are some agreements in the future that both agents prefer over the one that they reached. However, since A gains over time, it prefers to delay reaching an agreement as long as possible. Therefore, even if it promises W to give it a better offer in the future, when the time comes, it is not rational for A to keep its promise. A would rather wait another time period. Knowing that, W accepts a lower offer and an agreement is reached in the second period. So, even though the agreement is not Pareto optimal, the advantage of the solution is that an agreement is reached without a big delay, even though one of the agents gains over time. In that time period, there is no other agreement that is preferred by both agents.

- **Simplicity** — The strategies are simple and depend only on the current period. In the next section we will explain how these strategies can be implemented in an automated negotiator.
- **Symmetry** — The negotiation mechanism is sensitive to agent's roles (A or W), but not to their identity and types. The agents' type and identity only influence the other's private negotiation strategy.
- **Stability:** In all the situations that we considered we found sub-game perfect equilibrium or sequential equilibriums.
- **Satisfiability or accessibility:** In all the cases of task distribution that we considered the task will be performed. When we consider the resource allocation problem, there are some situations where an agent won't get the resource immediately. The first one is in the case of incomplete information, when one of the agents opts out. We assume that after the resource is repaired, both agents have some probability of getting access to the resource. The second case was described in Section 3. Here, if the agent's utility functions include a constant cost of delay, and $c_1 < c_2$ in the agreement that will be reached, the first agent will get all the resource. This is due to the structure of the situation, and giving up the usage of the resource is the second agent's preference.

9 Complexity and Implementation

We are in the process of implementing an automated negotiator that will participate in negotiation situations characterized by time constraints. It will be based on the theoretical results of this paper. We will report on the implementation issues in a different paper, but we discuss here some important related questions.

We have suggested that autonomous agents will use equilibrium strategies for the negotiation. Such strategies are stable since no designer will benefit by building an agent that uses any other strategy when it is known that the other agents are using equilibrium strategies.

There are two approaches to finding equilibria. One is the straight game theory approach: search for Nash or perfect or sequential strategies. The other is the economist's standard approach: set up a maximization problem and solve using calculus [43]. The maximization approach is straightforward and if the utility functions of the agents are chosen correctly, the maximization problem can be solved using some well known techniques of linear programming (e.g., [55]). However, when applied to situations such as ours, the maximization technique is less appropriate because the agents must solve their optimization problems jointly: A's strategy affects W's maximization problem and vice versa.

The drawback of the game theory approach is that finding equilibrium strategies is not mechanical: an agent must somehow make a guess that some strategy combination is in equilibrium before it tests it. There is no general way to make the initial guess.

In situations of multistage negotiations (or games in general) strategies can be found by trying to "guess" the set of actions that are used with positive probability in *each* state of the game. Working with this guess an agent can either construct a sequential equilibrium or show that none exists with this guess and go on and try another guess. If there is a point where it is clear that the negotiation will end it is often best to work through problems like this backward.

In our negotiation protocol there are $\frac{M^{\mathcal{A}-1}}{(|\mathcal{A}|-1)!}$ possible offers the offering agent can make (i.e., the number of possible agreements) and $3^{|\mathcal{A}|}$ possible combinations of actions with which the other agents can respond. If we assume that there is some time period $\hat{T}$ after which no agreement can be reached (e.g., W will prefer opting out to agreement), the overall number of pure strategies is $O((\frac{M^{|\mathcal{A}|-1}}{(|\mathcal{A}|-1)!})^{\hat{T}})$. However, if there is no such $\hat{T}$ then the number of pure strategies (in theory) can be infinite.

If we consider cases of incomplete information the problem becomes even harder. In addition to guesses of the actions, there should be a construction of the agents' beliefs in each state of the negotiation. This can be done by stating a set of inequalities which are the constraints on the beliefs in each state.

In general, it is too time consuming to compute the strategies in real time. Therefore, we suggest that equilibrium strategies be identified before the negotiation process starts. In this paper we presented appropriate strategies for varied situations. The situations are characterized by several environmental factors (e.g., number of agents, purpose of the negotiation) and the agents' utility functions.

If the automated agents act in a static environment, the appropriate strategy can be precomputed by the designer of the agent and inserted into its database. If the agents are active in a dynamic environment we are in the process of developing a library of frame-strategies that consists of the strategies that are appropriate for the different situations the automated negotiator may participate in. When the automated negotiator acts in one of these situations it will choose the frame-strategy that is appropriate for its current situation, initialize the appropriate variables (e.g., $\hat{s}^{W_l,t}$, c_A), and negotiate according to this strategy.

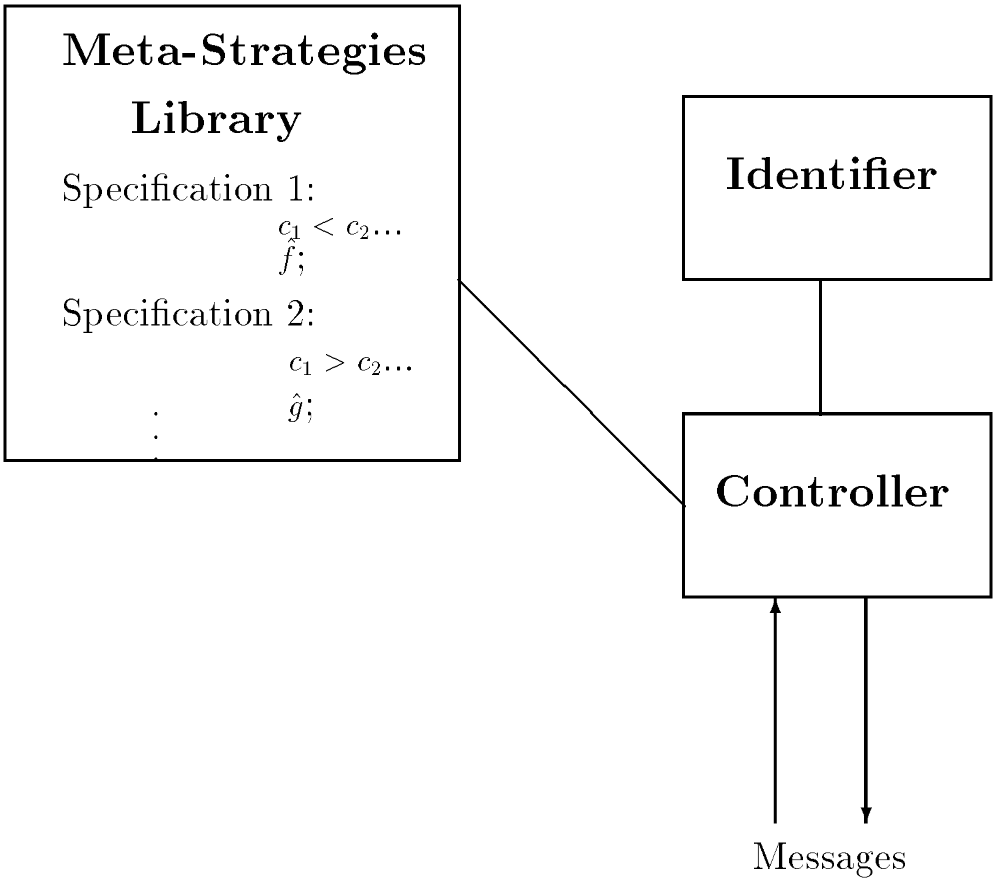

Figure 3: General Structure of the Automated Negotiator. The Identifier is responsible for determining the specific arguments of the situation. The Controller should choose a strategy from the Meta-Strategies' Library, according to the parameters. Then the Controller will operate according to the strategy and will send messages to other agents and receive messages from other agents.

In general, the automated negotiator will be composed of three modules: the Meta-Strategies Library, the Identifier and the Controller (see Figure 3).

A meta-strategy is a frame strategy that includes several parameters. We note that in each of the negotiation situations that we have investigated until now, the perfect equilibrium strategies are determined by parameters of the situation. For example, in the case in which the agents can opt out and they have constant delays (c_i), the strategies depend on the constant delays and the worst agreement for an agent which is still better for it than opting out in period 1 ($\hat{s}^{i,1}$). We use these parameters to construct the meta-strategies.

The Identifier will depend on the environment, and its purpose is to identify the special parameters of the environment (see [48]). The Controller will find the correct meta-strategy in the library, convert the meta-strategy into a strategy by initiating its parameters, and operate it.

So, in general, when the agent participates in any of the negotiation situations that we have considered, the Identifier will recognize the parameters of the situation (or they will be given to the agent), and the Controller will construct the exact strategy for the specific case and use it in the negotiations. Since we provide the agents with unique perfect equilibrium strategies, if we announce it to the other agents in the environment, the other agents cannot do better than to use their similar strategies. Since both the Meta-Strategies Library and the Controller are domain independent, the automated negotiator can be used in a variety of different applications, simply by changing the Identifier.

10 Conclusion

This paper has been concerned with how automated agents can be designed to interact effectively in both resource allocation and task distribution environments. A strategic model of negotiation has been proposed as a way of reaching mutual benefit while avoiding costly and time consuming interactions which might increase the overhead of coordination. That is, we have provided a model in which agents can avoid spending too much time negotiating an agreement and therefore are better able to stick to a timetable for satisfying their goals.

In the process of developing and specifying the strategic model of negotiation, we have examined single as well as multi-agent environments, situations characterized by complete as well as incomplete information, and the differing impact of time on the payoffs of the participants. While some combinations of these factors can result in minor delays, the model nevertheless reveals an important capacity for reaching agreement in early periods of the negotiation.

Throughout the paper, we have referred to two examples of application of the strategic model to problems in distributed artificial intelligence (DAI). The resource allocation problem has been examined through the development of a scenario in which agents must share a resource in order to achieve their separate goals. The task distribution problem has been examined via a scenario in which savings can result from the sharing of tasks, and both parties benefit from cooperation. In both cases, we have met the criteria for the evaluation of a negotiation protocol which we proposed at the outset of the paper: symmetrical distribution (no central unit or agent), efficiency (conflict avoided and no deadlocks in outcome), simplicity (process simple and efficient), stability (distinguishable equilibrium point), and satisfiability or accessibility (access to the resource or task completed).

We have ended with a brief discussion of the general structure of an automated negotiator, to be based on the theoretical results of this paper. The functioning of the automated negotiator will depend upon whether it will be operating in a static or dynamic environment. The implementation of a prototype automated negotiator will provide an environment in which experimental work on the strategic model under varying initial assumptions can be undertaken, as well as one in which human negotiators can be trained.

We believe that our model can be useful in other situations beside the ones we analyzed in the paper. For example, situations where there are several resources in the environment, or task distribution situations where the agents have incomplete information. We leave this for future work.

References

[1] A. H. Bond and L. Gasser. An analysis of problems and research in DAI. In A. H. Bond and L. Gasser, editors, *Readings in Distributed Artificial Intelligence*, pages 3–35. Morgan Kaufmann Publishers, Inc., San Mateo, California, 1988.

[2] S. Cammarata, D. McArthur, and R. Steeb. Strategies of cooperation in distributed problem solving. In *Proceedings of IJCAI-83*, 1983.

[3] N. Carver, Z. Cvetanovic, and V. Lesser. Sophisticated cooperation in FA/C distributed problem solving systems. In *Proc. of AAAI-91*, pages 191–198, California, 1991.

[4] B. Chandrasekn. Natural and social system metaphors for distributed problem solving: Introduction to the issue. *IEEE Transaction on Systems Man and Cybernetics*, 11(1):1–5, 1981.

[5] S. E. Conry, R. A. Meyer, and V. R. Lesser. Multistage negotiation in distributed planning. In A. H. Bond and L. Gasser, editors, *Readings in Distributed Artificial Intelligence*, pages 367–384. Morgan Kaufmann Publishers, Inc., San Mateo, California, 1988.

[6] S.E. Conry, , K.Kuwabara, V.R. Lesser, and R.A. Meyer. Multistage negotiation for distributed satisfaction. *IEEE Transactions on Systems, Man, and Cybernetics, Special Issue on Distributed Artificial Intelligence*, 21(6):1462–1477, December 1991.

[7] K. Decker and V. Lesser. A one-shot dynamic coordination algorithm for distributed sensor networks. In *Proc. of AAAI-93*, pages 210–216, Washington D.C., 1993.

[8] J. Doyle. Reasoning, representation, and rational self-government. In *Proc. of the 4th International Symposium on Methodologies for Intelligent Systems*, pages 367–380, 1989.

[9] J. Doyle. Rationality and its role in reasoning. In *Proc. of AAAI-90*, Boston, MA, 1990. Invited Talk.

[10] D. Druckman. *Negotiations*. Sage Publications, 1977.

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

[12] E. H. Durfee and V.R. Lesser. Global plans to coordinate distributed problem solvers. In *Proc. of IJCAI-87*, pages 875–883, 1987.

[13] E. Ephrati and J. Rosenschein. The clarke tax as a consensus mechanism among automated agents. In *Proc. of AAAI-91*, pages 173–178, California, 1991.

[14] R. E. Fikes, P.E. Hart, and N. J. Nilsson. Learning and executing generalized robot plans. *Artificial Intelligence*, 3(4):251–288, 1972.

[15] K. Fischer and N. Kuhn. A DAI approach to modeling the trasportation domain. Technical Report RR-93-25, Deustsches Forschungszentrum fur Kunstliche Intelligenz GmbH, 1993.

[16] R. Fisher and W. Ury. *Getting To Yes: Negotiating Agreement Without Giving In.* Houghton Mifflin, Boston, 1981.

[17] L. Gasser. Social knowledge and social action. In *IJCAI-93*, pages 751–757, French, 1993.

[18] M.R. Genesereth, M.L. Ginsberg, and J. Rosenschein. Cooperation without communication. In *Proc. of AAAI-86*, pages 51–57, Philadelphia, Pennsylvania, 1986.

[19] B.J. Grosz and C.L. Sidner. Plans for discourse. In P.R. Cohen, J.L. Morgan, and M.E. Pollack, editors, *Intentions in Communication*. Bradford Books at MIT Press, 1990.

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

[21] C.L Karrass. *The Negotiating Game: How to Get What You Want.* Thomas Crowell Company, NY, 1970.

[22] S. P. Ketchpel. Coalition formation among autonomous agents. In *Proc. of MAAMAW-93*, Neuchâtel, 1993.

[23] W. Kornfeld and C. Hewitt. The scientific community metaphor. *IEEE Transaction on Systems Man and Cybernetics*, 11(1):24–33, 1981.

[24] S. Kraus, E. Ephrati, and D. Lehmann. Negotiation in a non-cooperative environment. *J. of Experimental and Theoretical AI*, 3(4):255–282, 1991.

[25] S. Kraus and D. Lehmann. Designing and building a negotiating automated agent. *Computational Intelligence*, 1994. To appear.

[26] S. Kraus and J. Wilkenfeld. The function of time in cooperative negotiations. In *Proc. of AAAI-91*, pages 179–184, California, 1991.

[27] S. Kraus and J. Wilkenfeld. Negotiations over time in a multi agent environment: Preliminary report. In *Proc. of IJCAI-91*, pages 56–61, Australia, 1991.

[28] S. Kraus and J. Wilkenfeld. The updating of beliefs in negotiations under time constraints with uncertainty. In *Proc. of IJCAI-93 workshop on Artificial Economics*, pages 57–68, 1993. Also presented in BISFAI93.

[29] D. Kreps and R. Wilson. Sequential equilibria. *Econometrica*, 50:863–894, 1982.

[30] K. Kuwabara and V. Lesser. Extended protocol for multistage negotiation. In *Proc. of the Ninth Workshop on Distributed Artificiall Intelligence*, pages 129–161, Washington, 1989.

[31] V. R. Lesser and L. D. Erman. Distributed interpretation:a model and experiment. *IEEE Transaction on Computers*, 29(12):1144–1163, 1980.

[32] V. R. Lesser, J. Pavlin, and E. H. Durfee. Approximate processing in real time problem solving. *AI Magazine*, 9(1):49–61, 1988.

[33] V.R. Lesser. A retrospective view of fa/c distributed problem solving. *IEEE Transactions on Systems, Man, and Cybernetics, Special Issue on Distributed Artificial Intelligence*, 21(6):1347–1362, December 1991.

[34] K.E. Lochbaum, B.J. Grosz, and C.L. Sidner. Models of plans to support communication: An initial report. In *Proceedings of AAAI-90*, pages 485–490, Boston, MA, 1990.

[35] R. D. Luce and H. Raiffa. *Games and Decisions.* John Wiley and Sons, 1957.

[36] V. Madrigal, T. Tan, and R. Werlang. Support restrictions and sequential equililibria. *Journal of Economic Theory*, 43:329–334, 1987.

[37] T. W. Malone, R. E. Fikes, K. R. Grant, and M. T. Howard. Enterprise: A marketlike task schedule for distributed computing environments. In B. A. Huberman, editor, *The Ecology of Computation*, pages 177–205. North Holland, 1988.

[38] Y. Moses and M. Tennenholtz. Off-line reasoning for on-line efficiency. In *IJCAI-93*, pages 490–495, French, 1993.

[39] J. F. Nash. The bargaining problem. *Econometrica*, 18:155–162, 1950.

[40] J. F. Nash. Two-person cooperative games. *Econometrica*, 21:128–140, 1953.

[41] M. J. Osborne and A. Rubinstein. *Bargaining and Markets.* Academic Press Inc., San Diego, California, 1990.

[42] H. Raiffa. *The Art and Science of Negotiation.* Harvard University Press, 1982.

[43] E. Raizsmusen. *Games and Information.* Basil Blackwell Ltd., Cambridge, Ma, 1989.

[44] J. Rosenschein. *Rational Interaction: Cooperation Among Intelligent Agents.* PhD thesis, Stanford University, 1986.

[45] J. S. Rosenschein and G. Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiation Among Computers.* MIT Press, Boston, 1994. To appear.

[46] A. E. Roth. *Axiomatic Models of Bargaining.* Lecture Notes in Economics and Mathematical Systems No. 170. Spring-Verlag, Berlin, 1979.

[47] A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–109, 1982.

[48] K. Ruoff, D. de Hilster, A. Horry, L. Johnston, C. Kowalski, A. Meyers, and G. Vamos. Cooperation among heterogeneous intelligent subsystems in a military intelligence processing system. In *Proc. of the AAAI-91 workshop on Cooperation Among Heterogeneous Intelligent Agents*, California, 1991.

[49] T. Sandholm. An implementation of the contract net protocol based on marginal cost calculations. In *Proc. of AAAI-93*, pages 256–262, Washington D.C., 1993.

[50] R. Selten. Re-examination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory*, 4:25–55, 1975.

[51] A. Shaked and J. Sutton. Involuntary unemployment as a perfect equilibrium in a bargaining model. *Econometrica*, 52(6):1351–1364, 1984.

[52] O. Shechory and S. Kraus. Coalition formation among autonomous agents: Strategies and complexity. In *Fifth European Workshop on Modelling Autonomous Agents in a Multi-Agent World*, 1993.

[53] Y. Shoham and M. Tennenholtz. On the synthesis of useful social laws for artificial agent societies. In *Proc. of AAAI-92*, pages 276–281, San Jose, California, 1992.

[54] R.G. Smith and R. Davis. Negotiation as a metaphor for distributed problem solving. *Artificial Intelligence*, 20:63–109, 1983.

[55] W. Spivey and R. Thrall. *Linear Optimization.* Holt, Rinehart and Winston, INC, 1970.

[56] I. Stahl. An n-person bargaining game in an extensive form. In R. Henn and O. Moeschlin, editors, *Mathematical Economics and Game Theory*, Lecture Notes in Economics and Mathematical Systems No. 141. Berlin:Springer-Verlag, 1977.

[57] K. P. Sycara. Persuasive argumentation in negotiation. *Theory and Decision*, 28:203–242, 1990.

[58] K.P. Sycara. *Resolving Adversarial Conflicts: An Approach to Integrating Case-Based and Analytic Methods.* PhD thesis, School of Information and Computer Science, Georgia Institute of Technology, 1987.

[59] T. Tan and S. Werlang. A guide to knowledge and games. In *Proceedings of the Second Conference on Theoretical Aspects of Reasoning about Knowledge*, pages 163–177, Pacific Grove, California, March 1988.

[60] M. Wellman. A general-equilibrium approach to distributed transportation planning. In *Proc. of AAAI-92*, pages 282–289, San Jose, California, 1992.

[61] G. Zlotkin and J. Rosenschein. Cooperation and conflict resolution via negotiation among autonomous agents in noncooperative domains. *IEEE Transactions on Systems, Man, and Cybernetics, Special Issue on Distributed Artificial Intelligence*, 21(6):1317–1324, December 1991.

[62] G. Zlotkin and J. Rosenschein. One,two,many: Coalitions in multi-agent systems. In *Proc. of MAAMAW-93*, Neuchâtel, 1993.

[63] G. Zlotkin and J. S. Rosenschein. Negotiation and task sharing among autonomous agents in cooperative domains. In *Proceedings of IJCAI-89*, pages 912–917, Detroit, Michigan, August 1989.

[64] G. Zlotkin and J. S. Rosenschein. A domain theory for task oriented negotiation. In *Proceedings of IJCAI-93*, pages 416–422, Chambery, France, August 1993.

List of Figures