Baltzer Journals July 1,1996

Beliefs, Time and Incomplete Information in Multiple Encounter Negotiations Among Autonomous Agents*

SARIT KRAUS[1,2]

[1] *Department of Mathematics and Computer Science*
Bar Ilan University
Ramat Gan, 52900 Israel
E-mail: `sarit@cs.biu.ac.il`

[2] *Institute for Advanced Computer Studies,*
University of Maryland, College Park, 20742, MD

In negotiations among autonomous agents over resource allocation, beliefs about opponents, and about opponents' beliefs, become particularly important when there is incomplete information. This paper considers interactions among self-motivated, rational, and autonomous agents, each with its own utility function, and each seeking to maximize its expected utility. The paper expands upon previous work and focuses on incomplete information and multiple encounters among the agents. It presents a strategic model that takes into consideration the passage of time during the negotiation and also includes belief systems. The paper provides strategies for a wide range of situations. The framework satisfies the following criteria: symmetrical distribution, simplicity, instantaneously, efficiency and stability.

1 Introduction

Negotiation is an important mechanism for autonomous agents which have no central control and which need to reach an agreement on resource allocation. One of the main factors in such negotiations is the agents' beliefs about their opponents, and about their opponents' beliefs, and so on. Beliefs are especially important when the agents have only incomplete information about one another. Questions such as "what is the other agent's reservation price?," "how will my opponent respond if I reject his offer?," and "is it worthwhile for me to pretend

*This material is based upon work supported by the National Science Foundation under Grant No. IRI-9423967 and the Israeli Science Ministry grant No. 6288. I thank Jonathan Wilkenfeld and Rina Scwartz for helpful comments on earlier drafts and Onn Shehory for helpful discussions, particularly in proving Lemma 12.

to be someone else?" are common among negotiators. These questions become even more crucial when time is valuable and there are several possible agents with which to cooperate. Under these conditions, the agents may not be able to determine exactly what their opponents' beliefs are and therefore they will be unable to negotiate to the best of their capacity. In some situations, an agent needs to decide whom to negotiate with and how to estimate the possible results from negotiations with the other agents. We consider the interactions among *self-motivated*, *rational* and *autonomous* agents. We assume that each agent has its own utility function, and that rational behavior involves maximizing its expected utility. Our work belongs to the DAI class of *Multi-Agent Systems(MA)* (e.g., [52, 61, 63, 65]) rather than to the *Distributed Problem Solving (DPS)* class (e.g., [58, 17, 41]) as discussed in Section 2 below.

In previous work [31, 32, 34] we have developed a formal strategic model of negotiation that takes into consideration the passage of time during the negotiation process. In this paper, we extend this model to deal with *incomplete information* in *multiple encounters* when the agents negotiate on sharing a resource. There may be a need for the agents to share the resource due to limited resources (e.g., roads, bridges, clean air). In other situations resources are available but the agents may still mutually benefit from sharing a common resource since their use may be expensive (e.g., printers, satellites). There is, however, a conflict among the agents since all of them would like a larger share of the resource or a larger time period for using this resource.

One of the main characteristics of the situation is incomplete information. An agent which negotiates with another agent may have incomplete information about its opponent's type and may not be sure of how the opponent will evaluate an offer or how it might compare the offer with other options. While in our previous work we assumed that the negotiation occurs only once, here we assume that the agents may meet several times in situations where negotiation may be necessary. Thus, whereas in the previous work there was no long term consideration, in this paper future encounters play an important role.

In this paper strategies are presented that agents may use to influence their opponent's beliefs immediately, so they could benefit in future encounters. Furthermore, they may take actions designed to collect information about their opponents. Using this negotiation mechanism autonomous agents have *simple* and *stable* negotiation strategies that usually result in *efficient* agreements without delays.

In the rest of this section we will present the resource allocation problem that we consider and give criteria for evaluation of regulations for MA negotiation, and specify the assumptions we make. In Section 2 we discuss related work and in Section 3 we describe the negotiation model. In Section 3.4 we introduce modified results from our previous work concerning situations where agents negotiate only once. These results serve as a basis for the multiple encounters case that we consider in this paper. In Section 4 we present strategies for two agents that may negotiate on two occasions, and in Section 5 we discuss the general case of multiple encounters. In both sections we consider a wide range of situations. Section 6 discusses two possible extensions of the model and Section 7 discusses issues related to complexity and implementation of negotiating agents. In the last section we conclude by assessing the results using the criteria of Section 1.2.

1.1 The resource allocation problem

We consider the case where two agents share a resource that can only be used by one agent at a time. When the two agents face the problem of needing to use a resource simultaneously, they must reach an agreement on a schedule that divides the usage of the resource among the agents.[1] It may happen that these agents need to use the same resource simultaneously in the future, but they are not sure about it and they reach an agreement only on the current time period. If they meet again in the future, they start negotiating again on a new schedule, but each may use the information it collected on its opponent in the previous negotiation.

One example of a shared resource could be a communications satellite. Its high launch and maintenance cost, makes it necessary for a company to join with other companies, even rival ones, to gain access to something otherwise inaccessible. However, sharing a common resource requires a coordination mechanism that manages the resource.

A coordination mechanism can be designed to deal with either of the extremes: a static division of frequencies or time slots or an on-line negotiation mechanism that dynamically resolves the conflicts over the usage of the mutual resource. There are, however, on this spectrum, coordination mechanisms that generate agreements on long term global schedules (an hour, a day, ...). In this paper we consider repeated on-line negotiations between agents, (possibly from different companies), for the period of time these agents would want to use the same resource. We assume that the agents try to maximize their own utilities and are individually rational.

The elapsed time between when the resource is needed and the time the agents actually gain access has a cost to the agents, which depend on their internal states, such as their task load, their disk space, etc.

[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.

1.2 Evaluation criteria

The designers of agents should agree upon a protocol for negotiations. Since they do not have control over agents which may belong to other companies, this protocol should be accepted by all designers. Given this protocol, each agent will choose the best strategy for itself in a specific situation.

There are several criteria that should be taken into consideration when designers of agents consider possible protocols for negotiation on resource allocation in MA systems.

- **Distribution** — The decision making process should be distributed. The process should not be managed by a central unit or agent.
- **Instantaneously** — Conflict should be resolved without delay.
- **Conflict Avoidance** — Conflict should be avoided when possible.
- **Efficiency** — The resource is not in use only when there is no agent in the group that currently needs the resource.
- **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.
- **Symmetry** — The coordination mechanism should not treat agents differently because of non-relevant attributes. In the situations that we consider the relevant attributes are the agents' utility functions and their role in the encounter. All other attributes, such as agent's name or manufacture are not relevant. That is, symmetry implies that given a specific situation, the replacement of an agent with another which is identical in respect to the above attributes, will not change the outcome of the negotiation.
- **Stability** — There should be a distinguishable equilibrium point to the negotiation protocol (considered as a game). 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.[2] Being a "simple strategy" means that it is feasible to build it into an automated agent.

[2] 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.

As was shown in the economics literature [55, 8], the introduction of incomplete information into the model will result in some inefficiency. Also, since we consider complicated, repeated encounter negotiation situations, it is time consuming to find the appropriate strategies and therefore not efficient to do this on-line. However, the strategies that are presented in this paper are simple and can be easily implemented and provide on-line resource allocation.

1.3 Assumptions

The situations under consideration in this paper are characterized by the following assumptions.

1. **Bilateral Negotiation** — Even if there are several agents in the environment, in a given period of time no more than two agents need the same resource. 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. **Rationality** —The agents are rational; they try to maximize their utilities and behave according to their preferences. They use sequential equilibrium strategies[3].
3. **Commitments are Kept** — If an agreement is reached both sides will honor it.[4]
4. **No Long Term Commitments** — Each negotiation stands alone. An agent cannot commit itself to any future encounters. However, agents may use information obtained in one encounter in future encounters.
5. **Resource Division Possibilities** — We assume that the usage time of the resource can be divided in a discrete way.
6. **Agents' types** — There is a set of agents' types characterized by their capabilities. All agents know these types.
7. **Agents' Identity** — Agents can accurately identify each other.
8. **Negotiation Protocol** — Agents use the Alternative Offers model described in Section 3.[5]

[3] The concept of sequential equilibrium is defined in Section 5.

[4] This assumption is reasonable when agreements are implemented immediately. Otherwise, it may conflict with the Rationality assumption. Therefore, we assume that either the world does not change in the time between signing and implementing an agreement, or that there is some mechanism to enforce an agreement. However, the last assumption may contradict the *Distribution* requirement of the previous Section.

[5] This model imposes only minimal restrictions on the agents interactions as we explain in Section 3.

9. Common Beliefs Assumptions (1)-(8) are common belief.[6]

2 Related Work

Research in DAI is divided into two basic classes: *Distributed Problem Solving (DPS)* and *Multi-Agent Systems (MA)* [4]. Research in DPS considers how work involved in solving a particular problem can be divided among a number of modules or "nodes." The modules in a DPS system are designed to improve performance, stability, modularity, and reliability of the system. The modules include 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 intelligent agents. There is no global control, no globally consistent knowledge, and no globally shared goals or success criteria in MA. There is, however, a possibility for *real competition* among the agents. These are the two poles of the DAI research. Our research falls closer to the MA systems pole. We consider the problem of resource allocation in MA systems emphasizing the aspects of incomplete information, time constraints and multiple encounters. Other works in the DAI community dealing with the resource allocation problem (e.g., [12, 40, 47]) were closer to the DPS pole. In these works, as we discuss below, the problem of resource allocation arises from local conflicts among the agents, with each attempting as best as it can to fulfill its sub-tasks and contribute to the overall task of the system.

The issue of incomplete information adds an important dimension to the problem. Since the early 1980's, different models of sequential bargaining with incomplete information have been developed by economists and game theory researchers (e.g., [55, 8, 9, 2]), and as in our models of DAI situations in this paper and in [34], it was shown for the economics situations that the introduction of incomplete information tends to produce some inefficiency into the environment. The inefficiency can be either a delay in reaching an agreement (in our case agreements may be reached only in the second iteration as we explain in Section 3.4) or in a negotiation that ends without an agreement.

Another important issue is that of multiple encounters. Kreps and Wilson [35] and Milgrom and Robert [46] developed formal models that explain the common observation that in multi-stage "games", especially in industrial organizations, players may seek early in the game to acquire a reputation of being "tough" or "benevolent" or something else. Kreps and Wilson [35] studied a simple game of two players called entrant and monopolist. In this game they demonstrate the "reputation effect" where the players take actions that seem costly when viewed

[6] We assume common beliefs which is a much weaker notion than common knowledge, and unlike common knowledge [23], can be achieved in a distributed environment [29, 10].

in isolation but yield a reputation that is beneficial later. Milgrom and Robert [46] identified two factors that lead to the emergence of reputation: information uncertainty and repeated actions with the possibility of observing past behavior. They mention the choice of product quality and credit relationships among the situations where these factors play an important role. We also consider situations of incomplete information and repeated actions (encounters), and thus reputation emerges in our cases too.

The details of the situations we consider are quite different from those of [35, 46]. In each of our encounters the agents use the model of alternative offers to reach an agreement on resource division; thus our model is a combination of multi-stage "games" where each stage is composed of multiple encounters. However, we are able to use similar techniques of sequential equilibrium which we present in Section 5.

Another related model is the repeated sale model [26, 5, 6] where the same agents negotiate several times. These situations are similar to ours, but while they consider cases of buyer/seller paradigms or landlord/lessee paradigms, we examined situations of resource allocation among multiple agents. In our case there are only short term agreements. We apply game theory techniques to scenarios that were not considered by the game theory researchers.

The issue of incomplete information in DAI was studied mainly in DPS environments. In Davis and Smith's work on the Contract Net [58], they present nodes which have incomplete information about the other nodes' load and their possibility to carry out sub-tasks. Davis and Smith developed a form of simple negotiation among the cooperative nodes, with one node announcing the availability of tasks and awarding them to other honest bidding nodes. The bidding nodes do not try to manipulate the situation or to transfer misleading information since all the nodes are working on the same task.

Malone et al. [45] developed a Distributed Scheduling Protocol (DSP) based on the contract net protocol for their Enterprise system. The most important way in which DSP differs from the original contract net protocol is by its criteria for matching between tasks and agents (i.e., the problem of sub-tasks distribution). It includes two primary dimensions: (1) contractors select managers' tasks in the order of the tasks' numerical *priorities*, and (2) managers select contractors on the basis of *estimated completion times* from among the contractors that satisfy the minimum requirements to perform the job. Malone et al. also considered the problem of dis-information. Since they allow people to supply their own estimation of processing times for their tasks and these time estimates are also used to determine priority, there is a clear incentive for people to bias their processing time estimates in order to get higher priority. However, in Enterprise, the node performing a task (i.e., the contractor) knows the correct time of the performance while carrying the task, and can use it for imposing sanctions on the clients. If a task takes significantly longer than it was estimated to take, the contractor aborts the task and notifies the client that it was "cut off". This cutoff feature prevents the possibility of a few people or tasks monopolizing the system. This technique is not useful in our framework since after the distribution of the resource, there is no way for the other agent to gain more information on its opponent. Therefore, there is no way to verify whether an agent tells the truth or not when announcing its preferences. There is no usage for such announcements. The only source of information about an agent is its actions.

A modified version of the Contract Net protocol for competitive agents in the transportation domain is presented in [56]. It provides a formalization of the bidding and the decision awarding processes, based on marginal cost calculations of local agent criteria. More important, an agent will submit a bid for a set of delivery tasks[7] only if the maximum price mentioned in the tasks' announcement is greater than what the deliveries will cost that agent. A simple motivation technique is presented to convince agents to make bids; the actual price of a contract is half way between the price mentioned in the task announcement and the bid price. Sandholm considers task negotiation rather than negotiation over resources as we do. In this context he presents heuristics to problems we don't consider, such as, how to choose which tasks to contract out, how to cluster tasks into sets to bargain over as atomic bargaining and how to bid when multiple bids and awards should be handled simultaneously. On the other hand, there is no discussion of how manipulation of the task announcements can affect the behavior of the system and bidding and awarding decisions do not anticipate future contracts. Also, the time of the negotiation is not taken in to consideration. We concentrate on these aspects of negotiation in the context of resource allocation.

Lesser and his colleagues [7, 47, 38, 13] considered the problem that agents working as a team may possibly form different views of the situation. They therefore suggest different frameworks for negotiation and communications for information exchange and conflict resolution. Since the agents are cooperative they are assumed to be honest. For example, Conry at el. and Kuwabara and Lesser [11, 37, 12] presented a multistage negotiation protocol that is useful for cooperatively resolving resource allocation conflicts arising in distributed networks of semi-autonomous problem solving nodes. Since they consider the case of scheduling of many resources for multiple tasks they allow agents to negotiate with different agents simultaneously. We concentrate on negotiation between two agents emphasizing the issue of the negotiation time and providing fast negotiation strategies.[8]

[7] Announcing one delivery at a time is not sufficient in general. This is due to the fact that the deliveries are dependent. For example, for two disjointed delivery sets T_1 and T_2, the marginal costs that are saved by removing both T_1 and T_2 are usually larger than the sum of marginal cost that was saved by removing each of them alone.

[8] Even though in the paper we consider the case of the division of one resource, it is easy to

[47] describes a framework called DENEGOT for negotiating conflicts that arise in multi-agent planning with time and cost constraints. Top-level goals are originally predefined with some threshold level of global cost and time utility required. Agents own resources and have predefined responsibilities. If an agent cannot find a plan to optimally meet the top-level goals it is responsible for with its own resources, the agent negotiates with other agents to borrow resources to help achieve its goals. The intent of the negotiation is to find a combined multi-agent plan in which all the top-level goals are satisfied in an acceptable, though not necessarily optimal, fashion. The agents view this negotiation process as a distributed search and the main purpose of the negotiation is to exchange information among cooperative agents, rather than reaching an agreement among opponents as in our framework.

Zlotkin and Rosenschein [66, 67, 53] studied the problem of incomplete information in negotiation in MA systems where the agents need to reach an agreement on task allocation. The incomplete information is either about the opponent's goals or about the value of its goals. They introduce a mechanism that they called "-1 negotiation phase" in which agents simultaneously declare private information before beginning the negotiation. Zlotkin and Rosenschein also identified situations and protocols where agents have incentives to tell the truth in "-1 negotiation phase" and cases where it is beneficial for the agents to lie. In our model there is no pre-negotiation phase. Information gathering can be done only during the negotiation by the agents' proposals and actions. Zlotkin and Rosenschein assume that the agents negotiate only once, while we consider multi-encounter situations. Additionally they consider the task distribution problem while we consider the resource allocation problem. Another key difference between our model and theirs is that our model takes into consideration the passage of time during the negotiation process itself, which in turn influences the outcome of the negotiations and avoids delays in reaching an agreement. The passage of time is not considered in Zlotkin and Rosenschein's work.

Ephrati and Rosenschein ([18, 19]) used the Clarke Tax voting procedure as a consensus mechanism when there was incomplete information about the values of agents' goals. The mechanism assumes an explicit utility transferability (i.e., a kind of monetary system) which is not available in our framework. In addition they considered the case of only one encounter between the agents.

Sycara [62] presented a model of negotiation that combines case-based reasoning and optimization of the agents' multi-attribute utilities. She implemented her ideas in a computer program called the PERSUADER which resolved conflicts in the domain of labor relations, and tested her system using simulations of such domains. In her system, agents' beliefs about other agents' beliefs and

extend the model to the case of two agents reaching an agreement over the division of several resources.

goals change during the negotiation. While she concentrated on the perspective of the mediator in single encounter negotiations, we consider the negotiation process from the point of view of the automated negotiators in multiple encounters.

In [30], *Diplomat*, an automated agent that plays Diplomacy was designed and implemented. Diplomacy is a game of incomplete information with multiple encounters. Diplomacy players have incomplete information about their opponents' goals and tasks and about the coalitions that were formed between the other players. Under these circumstances, agents may tell lies and may not keep their promises. One of Diplomat's main efforts is to try to estimate what its opponent's goals are and whether they will keep their promises. It also tries to mislead its opponents in order to increase its own benefits, but it will also try to maintain its reputation and credibility for future encounters. Whereas in [30] the agents use heuristics to reach these goals, in this paper we provide a formal model and find equilibrium strategies.

In our work agents revise their beliefs about their opponents during the negotiation. The question of how an agent should revise its beliefs has long occupied philosophers of science and statisticians (e.g., [57, 25]). Knowledge is often viewed in probabilistic terms, thus revising beliefs becomes identical to updating probabilities over assertions which, in some sense is similar to our approach.[9]

Dubois and Prade [16] provide a survey of revision and updating operations available in probability theory and in the possibility theory framework. They examine the two main ways that are offered by the probabilistic framework to modify a probability distribution upon the arrival of new information: the Bayesian conditioning and the 'imaging' which consists of translating the weights originally on worlds outside a given world A, toward worlds which are their closest neighbors in A. They show that these two modes are analogous to the distinction between *belief revision* based on Alchourrón, Gärdenfors and Makinson's postulates [22] and *updating* based on Katsuno and Mendelzon's postulates [28]. Our techniques belong to the "revision" paradigm. We have examined situations where an agent has initial probabilistic beliefs which are revised using Bayesian rules, when the agent observes the actions of its opponents in negotiation. The revision is done in the context of the equilibrium strategies and is based on the actions the opponents are supposed to take according to these strategies.

[9] Belief revision has more recently been treated by philosophers (e.g., [49, 42, 60, 24, 22]), theoretical computer scientists (e.g., [20]), and artificial intelligence researchers (e.g., [14, 64, 50, 27, 15]). These groups view an agent's beliefs as a set of assertions (without probabilities), and revising beliefs involves deciding how that set of assertions should change when new information arrives.

3 The Negotiation Protocol

Our strategic model of negotiation is a modification of Rubinstein Alternative Offers model [54, 55]. We utilize modified definitions from [34].[10].

We assume here that there is a set of agents denoted by **Agent**. The negotiation is between two agents that negotiate the division of M units of a resource.

Definition 1 (agreement)
An agreement *is an ordered pair* (s_1, s_2)*, where* $s_i \in \mathbb{N}$ *and* $s_1 + s_2 = M$*.* s_i *is agent* i*'s portion of the resource or task. We denote by* S *the set of all possible agreements.*

Negotiation is a process that may include several iterations and may even continue forever. Each iteration takes two steps. In the first step of any negotiation iteration, one agent[11], say i, proposes an agreement from S. In the next step, the other agent (j) either accepts the offer (**Yes**) or rejects it (**No**), or opts out of the negotiation (chooses **Opt**). 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). If j chooses opting out, the negotiation also ends. 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 very general framework for the negotiation process and specifies that agents should respond to offers and make counteroffers. The framework specifies termination conditions, but there is no limit to the number of iterations[12]. We denote the set of negotiation iterations and call each a "time period" by the ordered set $\mathcal{T} = \{0, 1, 2, ...\}$.

As mentioned in Section 1.3 we will assume that there is a finite set of agent types characterized by their capabilities (e.g., their disk space, computational power, payment agreements). These characteristics produce a different utility function for each type of agent. Assuming that each agent i has a utility function over all possible outcomes: $U^i\colon \{\{S \cup \{\mathbf{Opt}\}\} \times \mathcal{T}\} \cup \{\mathbf{Disagreement}\} \rightarrow \mathbb{R}$. In addition, each agent has some probabilistic beliefs about the types of the other agents, and about the other agents' beliefs about themselves and about other agents. These beliefs may be updated over time, during negotiations between the agents. Formally, we denote by $\mathbf{Type} = \{1, ..., k\}$ the possible types of the agents. We assume that the details of those types are mutually believed by the agents.

[10] See [48] for a detailed review of the bargaining game of Alternative Offers.

[11] We assume that the agent that needs the resource will start the negotiation.

[12] In previous work we assumed that an offer and the response occur in the same negotiation step. We make this change to enable the agents to update their beliefs after receiving an offer. When there is complete information, the models are equivalent.

An agent's negotiation strategy is, in general, any function of the history of the negotiations to its next move. In order to formally define a strategy we will define the notions of history and of an agent's belief.

Definition 2 (history)
For any time period $t \in \mathcal{T}$ of the negotiation let $H(t)$ be the history through time period t of the negotiation. $H(t)$ is a sequence of $\lfloor \frac{t}{2} \rfloor + 1$ proposals and $\lfloor \frac{t}{2} \rfloor$ responses.

For example, suppose there are two agents and $M = 50$. If in the first time period agent i proposes $(30, 20)$ which was rejected by agent j then $H(1) = \{(30, 20)\}$ and $H(2) = \{(30, 20), \mathbf{No}\}$. If in the third time period agent j proposes $(25, 25)$ which is accepted by agent i then $H(3) = \{(30, 20), \mathbf{No}, (25, 25)\}$ and $H(4) = \{(30, 20), \mathbf{No}, (25, 25), \mathbf{Yes}\}$.

Definition 3 (system of beliefs)
*A system of beliefs of agent i is a function $\wp_i(H)$ which is a probability distribution of i's opponents as a function of the history. That is, $\wp_i(H) = \{(\phi_1^j, ..., \phi_k^j) | j \in$ **Agent** $\setminus \{i\}\}$ describes agent i's belief about its opponents' types according to a given history of offers and counteroffers H.*

For example, suppose there are two agents i and j and three types of agents in the environment, and suppose that before the negotiation starts agent i 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, $\wp_i(\emptyset) = \{(\frac{1}{2}, \frac{1}{4}, \frac{1}{4})\}$. Now suppose i receives an offer s from its opponent j. i may now change its beliefs. For example, it may conclude that its opponent cannot be of type 3, but rather there is probability $\frac{2}{3}$ that it is of type 1 and probability $\frac{1}{3}$ that it is of type 2. That is, $\wp_i(\{s\}) = \{(\frac{2}{3}, \frac{1}{3}, 0)\}$.

Using these definitions, we will describe the notions of *pure* and *mixed* strategies that were proposed by Von Neumann [43]. As mentioned above, a pure strategy specifies an action for an agent which can be either a proposal or a response, given the history and the system of beliefs. A mixed strategy requires the agent to draw a random number with a probability distribution specified by the strategy, and then decide accordingly on the action it will take. These mixed strategies will be used to find stable solutions in situations where there are no stable pure strategies.

Definition 4 (strategies)
A pure strategy *for an agent i specifies an action in the set* $\{$**Yes, No, Opt**$\} \cup S$ *for every system of beliefs and possible history after which this agent has to take*

an action.

A mixed strategy *for an agent specifies a probability distribution over actions rather than just an action as in the pure strategies.*

3.1 Sequential equilibrium

The main questions here pertain to how an agent uses its beliefs during the negotiation, how it updates its beliefs according to the information it gathers during the negotiation process, and how an agent influences its opponents' beliefs. We examine these problems in several situations using the notion of *sequential equilibrium* [36], which requires that in each time period any agent's strategy will be optimal given its opponents' strategies, the history up to the given time period, and its beliefs. The agent's beliefs may change over time, but must remain consistent with the history.

In order to state the requirement that an agent's strategy be optimal for every history, we must specify its beliefs about the other agents' types. The notion of sequential equilibrium therefore, requires the specification of two elements: the profile of strategies and the beliefs of the agents. This means that, when the number of agents is n and the number of possible types is k then a **Sequential Equilibrium (S.E.)** is a sequence of nk strategies (i.e., k strategies for each agent for any possible type, $1_1, 1_2, .., 1_k, ..., n_1, ..., n_k$) and a system of belief with the following properties: each agent has a belief about its opponents' 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 is consistent with the history of the negotiation. Meaning, the agents' belief may change over time, but must remain consistent with the history. We assume that each agent in a negotiation interaction has an initial system of beliefs. While the agents' beliefs may change over time, its type, which is characterize by its capabilities and goals doesn't change over time as we explain below.

A sequence of nk strategies, one for each possible agent leads, from the point of view of the agents, to a probability distribution over outcomes. For example, if agent i believes with probability ϕ that its opponent j is of type 2, then i expects that with probability ϕ the outcome is determined by the strategy specified to i and the strategy specified in the sequential equilibrium to j_2. If i believes that j's type is k with probability ϕ', then it assumes that with probability ϕ' that the outcome will be the result of j's usage of the strategy that is specified in the sequential equilibrium for type k and its own strategy. The agents use expected utilities to compare among these outcomes.

We impose three conditions on the sequence of strategies and the agent's system of beliefs.[13]

- **Sequential Rationality** — The optimality of agent i's strategy after any history H depends on the strategies of its opponents, given their types and its system of beliefs. This means that agent i will try to maximize its expected utility, with regard to the strategies of its opponents and its beliefs about the probabilities of its opponents' type according to the given history.

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

  If, after any history, all strategies of agent j's, regardless of agent j's type, indicate that it has to take the same action (e.g., reject an offer, make the same counteroffer), and this action is indeed made by agent j, then agent i's beliefs remain the same as they were before the action was made. If only one of the strategies of j, for example, type l, specifies that a given action should be taken (e.g., making an offer s), and the action is indeed taken (e.g., s is indeed offered by j), then i believes with probability 1 that j's type is indeed l. The agent uses the same reasoning about its opponents beliefs and updates it in a similar way.

  To demonstrate this requirement we return to the above example, where there are three types of agents in the environment. Suppose i's original belief was $\wp_i(\emptyset) = \{(\frac{1}{2}, \frac{1}{4}, \frac{1}{4})\}$ as above, and suppose that the strategies of j_1, j_2 and j_3 indicate that in the beginning all of them will make an offer s, then i's beliefs cannot be changed if it indeed receives the offer s. However, if the strategies of j_1 and j_2 specify the offer s, but the strategy of j specifies the offer s', then if A receives an offer s' it believes that its opponent is of type 3. That is, $\wp_i(\{s\}) = \{(0, 0, 1)\}$.

- **Never Dissuaded Once Convinced** — Once an agent is convinced of its opponent's type with probability 1, or convinced that its opponent cannot 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 i reaches the conclusion that its opponent is j_3, it cannot revise its belief, even if agent j subsequently deviates from j_3's strategy. From this point on, i has perfect information on agent j_3 and it is sure how j will respond to its offers and which counteroffers it will make.[14]

[13] Kreps and Wilson [36] imposed an additional stronger restriction. They required that the beliefs of the agent are the limit of a sequence of rational beliefs. All the original sequential equilibria satisfied our conditions, but there are a few equilibria according to our definition that do not satisfy Kreps and Wilson's stronger requirement.

[14] See [44] for a discussion of this requirement. This requirement may cause, in some situations, the elimination of equilibria. We leave the relaxation of this requirement for future work.

Definition 5 (sequential equilibrium)
A sequential equilibrium *is a sequence of nk strategies and a system of beliefs, for any $i \in$* **Type** *that satisfy the conditions of* **Sequential Rationality**, **Consistency** *and* **Never Dissuaded Once Convinced**.

Using this formal definition of sequential equilibrium and the negotiation protocol, we will analyze different negotiation situations. The concepts of *pooling* and *separating* equilibria are very useful in analyzing situations of multiple encounters and reputation [21]. Suppose there are several types of agents. If all types of agents pick the same strategy in all states, the equilibrium is *pooling*. Otherwise, it is *separating*. There can also be *hybrid* or *semi-separating* equilibria where an agent may randomize between pooling and separating. We use these concepts later in the paper.

3.2 Probabilistic actions in multi-agents environments

In some of the situations that we consider, there is no sequential equilibrium with pure strategies. Therefore, we propose that the agents will use mixed strategies, i.e., they will randomly choose what to do next, according to the probabilities specified by the strategy. When the agents choose to randomize between several pure strategies, the expected utility from all of the chosen pure strategies should be the same. Otherwise, they will not agree to randomize but will prefer one pure strategy over the other.

Mixed strategies that are in sequential equilibrium are not as intuitive as pure strategies equilibrium, and many game theorists and economists prefer to restrict themselves to pure strategies in games that have both pure and mixed strategies equilibrium. Similarly, we suggest using mixed strategies only when there is no equilibrium with pure strategies. However, we claim that using mixed strategies for automated agents is a good technique.

Game theorists and economists try to model and estimate human behavior [51]. One of their main objections to mixed strategies is, that people in the real world do not take random actions. This observation is not applicable in MA systems where all agents are automated and the designers of agents can come to a general agreement that their agents use mixed strategies.[15] Even in the case where some of the agents are human, the automated agent can treat the mixed strategies as good descriptions of human behavior in the sense that the actions appear as random to observers, even if the human agent himself/herself has always been sure what action he/she would take. For example, if there are several types of human agents, each takes a different action, and the automated agent has some

[15] Note that this does not require central design of agents. However, it requires the development of some standard for MA environment.

probabilistic beliefs about the human's type. Moreover, explicitly random actions are not uncommon among humans. For example, the IRS's heuristics for deciding which tax return to audit include random actions.

Another objection to the usage of mixed strategies is that an agent which selects mixed strategies must always be indifferent between two pure strategies. Even a small deviation from the probabilities specified by the equilibrium destroys the equilibrium, while this deviation does not change the agents' expected utility. That is, to maintain the equilibrium, a player must pick a particular strategy from strategies it is indifferent between. It seems that in the case of automated agents, the designers can agree in advance on such behavior.

Zlotkin and Rosenschein also consider some sort of probabilistic actions. In [65] they proposed the notion of *mixed deals* in order to resolve conflicts in task distribution. Zlotkin and Rosenschein defined a mixed deal to be a pair of plans P_A and P_B and a probability p. If the agents reach this deal, then with probability p, agent A will do P_A and agent B will carry P_B, and with probability $1-p$, A will do P_B, and B will carry out P_A. That is, Zlotkin and Rosenschein's protocol requires that the agents need to draw the random number jointly. The expected utility of an agent from P_A and P_B is different and there should be some mechanism to force them to carry out their promises after they jointly draw the random number. [16]

Note that Zlotkin and Rosenschein's concept is very different from ours.[17] We propose to use only pure deals. An agent chooses a strategy randomly, in private, and is motivated by the property that the expected utilities of the strategies it mixes between them are the same. Furthermore, using Zlotkin and Rosenschein's mixed deals won't provide stability in our case. If an agent will agree on a mixed deal, it will reveal its type. This is not acceptable in the cases where it considers mixed strategies.

3.3 Attributes of the utility functions

In the rest of the paper we assume that there is one agent that is currently using a resource (A symbolizing "access"), and that there is another agent which also wants to use it (W symbolizing "waiting"). W wishes to gain access to the resource during the next M time periods.

First we modify several assumptions of [31, 32, 34] to fit a situation in which agents may be of several types. For $i \in \mathbf{Type}$, if the type of agent A (respectively W) is i, we denote it by A_i (respectively W_i). If a condition holds regardless of the type of agent A (respectively W) we use A (respectively W).

The first assumption states that agents prefer any agreement in any given

[16]This concept is similar to the notion of correlated equilibrium [1] which in most situations requires a contract enforcement mechanism.

[17]Similarly, correlated equilibrium is different from mixed strategies.

time period over the continuation of the negotiation process indefinitely.

$A0$ **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.

The next assumption requires that among agreements reached in the same period, agent i prefers larger portions of the resource.

$A1$ **The Resource is Valuable:** For all $t \in \mathcal{T}, r, s \in S$ and $i \in \mathbf{Agent}$: $r_i > s_i \Rightarrow U^i((r,t)) > U^i((s,t))$.[18]

The next assumption expresses the agents' different attitudes toward time. W is losing over time while A is gaining over time.

$A2$ **Cost Over Time:** For any $t_1, t_2 \in \mathcal{T}, s \in S$ and if $t_1 < t_2$, $U^W((s,t_1)) > U^W((s,t_2))$ and $U^A((s,t_1)) < U^A((s,t_2))$.

We assume that the agents have a utility function with a constant cost or gain due to delay. Every agent bears a fixed cost for each period. That is, each agent A_i has a constant time gain $c_{A_i} > 0$, and each agent W_i has a constant time loss, $c_{W_i} < 0$.

$A3$ **Agreement's Cost Over Time:** Each agent $i \in \{W_1, W_2, \cdots, W_k, A_1, ..., A_k\}$ has a number c_i such that: $U^i(s,t_1) \geq U^i(\bar{s},t_2)$ iff $(s_i + c_i\lfloor\frac{t_1}{2}\rfloor) \geq (\bar{s}_i + c_i\lfloor\frac{t_2}{2}\rfloor)$,[19] where for any $j \in \mathbf{Type}$, $c_{W_j} < 0$ and $c_{A_j} > 0$ and $|c_{W_k}| \leq |c_{W_{k-1}}| \leq \cdots \leq |c_{W_1}| \leq |c_{A_k}| \leq \cdots \leq |c_{A_1}|$.

That is, an agent will gain more while using the resource than it will lose while waiting for the resource.

The next assumption concerns the utility of opting out. W prefers opting out sooner rather than later (regardless of its type) and A always prefers opting out later rather than sooner (regardless of its type). This is because A gains over time while W loses over time. For this reason A would *never* opt out. A would prefer for agent W to opt out in the next iteration than to opt out by itself in the current iteration.

$A4$ **Cost of Opting Out Over Time:** for any $t \in \mathcal{T}$,
$U^W((\mathbf{Opt},t)) > U^W((\mathbf{Opt},t+1))$ and
$U^A((\mathbf{Opt},t)) < U^A((\mathbf{Opt},t+1))$.

[18] For all $s \in S$ and $i \in \mathbf{Agent}$, s_i is agent i's portion of the resource.

[19] t_1 and t_2 are divided by 2 to make the model similar to our previous one [32, 33, 34] where each iteration took only one time period.

Even though agent A prefers to continue the negotiation indefinitely, an agreement will be reached (after a finite number of periods) if the next assumption holds. 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.

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 more preferable to i than opting out in time period t. We denote this agreement by $\hat{s}^{i,t}$.

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. We also assume that 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 in the next iteration.

A5 **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 $\hat{s}_A^{W,t} \geq 0$ then $U^A((\hat{s}^{W,t}, t)) > U^A((\mathbf{Opt}, t+1))$ and $U^A((\hat{s}^{W,t+1}, t+1)) > U^A((\hat{s}^{W,t}, t))$. In addition, $\hat{s}_W^{W_1,t} < \hat{s}_W^{W_2,t} < \cdots < \hat{s}_W^{W_k,t}$.

Our last assumption is that in the first period there is an agreement that is preferable to both agents (regardless of their type) over opting out.

A6 **Possible Agreement:** For all $i, j \in$ **Agent**, $U^i((\hat{s}^{j,0}, 0)) \geq U^i((\hat{s}^{i,0}, 0))$

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

3.4 Agents' negotiation in a single encounter

We will first consider the case where the probability that the agents will meet again is very low, and therefore future encounters are not taken into consideration. The results of the single encounter case will be used when considering the multiple encounter situation.

Hence we recall from [34] what will happen in such cases (see detailed discussion and proofs in [34].[20]) Suppose $i, j \in$ **Type**.

1. In general, W_i will not offer A_j anything better than its possible utility from opting out in the next iteration, with the addition of W_i's loss over time (i.e., $\hat{s}_A^{W_i,t+3} + |c_{W_i}|$), since it can always wait until the next iteration and opt out.

[20] Note, that we adjust the results presented in this section to fit our new model, i.e., each iteration takes two time periods rather than one. However, since the loss over time depends on the number of iterations, the models are very similar.

2. If A_j receives such an offer, it may realize that its opponent is no stronger than type i. This is because a stronger agent will not offer $\hat{s}_A^{W_i,t+3} + |c_{W_i}|$. If it realizes that its opponent is at most of type i, it can wait until the next iteration and offer W $\hat{s}^{W_i,t+3}$. Since A_j gains more than W loses $|c_{W_i}| < c_{A_j}$, A_j will prefer it over W's current offer.

3. It is better for W_i to "pretend"[21] to be the strongest type by offering only $\hat{s}_A^{W_k,t+3} + |c_{W_k}|$. This offer will be rejected, but in this way it will not convey any information about W's type to A_j.

4. In the next iteration, if W_i gets an offer that is worth less than opting out (less than $\hat{s}^{W_i,t+3}$), then it will really opt out. That is, if A_j offers $\hat{s}^{W_l,t+3}$ and W is of a stronger type than l it will opt out. This is because W_i knows that in any future time period t', it will not reach an agreement better than $\hat{s}^{W_i,t+3}$, and W_i would prefer to opt out over that possibility. Therefore, A_j computes its expected utility for offering $\hat{s}^{W_i,t+3}$, for any $i \in \mathbf{Type}$, according to its beliefs, and offers the one where its expected utility is maximal. After receiving the offer, W will either accept it or opt out according to its type.

This means that, if the agents use sequential equilibrium strategies, the negotiation will end in the second iteration and 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_j's initial belief. Note that in this case, the agent's behavior is not influenced by its beliefs about its opponent's beliefs. Only each agent's own beliefs play a role here. However, A_j may update its belief about W's type at the end of the negotiation.

According to assumption $A5$, agent A always prefers $\hat{s}^{W,t}$ over opting out, regardless of its own type and regardless of W's type. Therefore, in most of the discussion on resource allocation, we will consider cases in which A_j's type does not play an important role, meaning that A_j's actions do not depend on its exact type. This simplifies our discussion.[22] Furthermore, in the next two sections we will assume that there are only two types of agents in the environment, "high" and "low", i.e., $\mathbf{Type} = \{h, l\}$ such that, $|c_{W_h}| \leq |c_{W_l}| \leq c_{A_h} \leq c_{A_l}$ and $\hat{s}_W^{W_h,t} > \hat{s}_W^{W_l,t}$. This means that agent h would prefer to opt out more often than agent l. Then the threat of h to opt out in case it gets low offers causes its opponent to offer it higher agreements.

[21] When we say that an agent B pretends to be agent C in a given situation, we mean that agent B will take the same action as agent C, regardless of its expected utility in this situation. However, if two agents are expected to take the same action, their opponent won't change its beliefs observing the common action, as was stated in the **Consistency** requirement of S.E.

[22] From a practical point of view, these are situations of asymmetric information.

3.5 Multiple-encounter negotiation protocol

The main concern in this paper is situations where two agents may negotiate several times. We will adjust the negotiation protocol described in the beginning of Section 3 to the multiple-encounter negotiation situation.[23]

The notion of *history* defined in Definition 2 describes the progress of the negotiation in a single encounter negotiation. In multiple encounters, given two agents, we use a sequence of histories to describe the encounters between them over time. For $i, j \in$ **Agent**, if the agents negotiate m times, we denote by $\mathcal{H}_{i.j} = H_1, ..., H_m$ the sequence of their histories. Furthermore, we assume that interactions with other agents will not change the agents' beliefs about one another.[24] Therefore, if there are several encounters, we assume that the beliefs of the agent at the beginning of encounter H_q $1 < q \leq m$ will be the same as at the end of the encounter H_{q-1}.

A *pure strategy* for an agent i specifies an action in the set $\{\mathbf{Yes},\ \mathbf{No},\ \mathbf{Opt}\} \cup S$ for every possible sequence of histories and appropriate system of beliefs. Since the agents' belief at the end of one history is similar to the one at the beginning of the next history in the sequence, there is no need for a strategy to be a function of all the histories in a sequence of histories. Therefore, a strategy for a sequence of histories will be composed of strategies that are functions of one history in the sequence.

Furthermore, since in this paper we concentrate on the effect of multiple encounters, we will not describe in detail the agent's strategies for each history in the sequence. We will use the strategies described in Section 3.4 as the basic components of strategies of sequences of histories. For strategies that form a sequential equilibrium, we will identify them with the actual events that occur. For example, given a specific encounter, when saying that A_j will offer $\hat{s}^{W_h,3}$, we mean that in every time period when it is A_j's turn to make an offer, it will offer $\hat{s}^{W_h,t}$ and if it receives an offer smaller than $\hat{s}^{W_h,t} + c_{A_j}$ it will reject it. However, since by the second iteration, given W's strategy, either agreement with $\hat{s}^{W_h,3}$ or $\hat{s}^{W_l,3}$ will be reached, or W will opt out, we will characterize the strategies by the behavior in the second iteration, i.e., $\hat{s}^{W_l,3}$, $\hat{s}^{W_h,3}$ or **Opt**. Thus, in the rest of the paper the main factors that play a role are the agents' utilities in the third time period. To make the paper more readable, we will use short notations for these utilities, U_j^i which denotes the utility of agent i from outcome j. In particular, we use l for $(\hat{s}^{W_l,3}, 3)$, h for $(\hat{s}^{W_h,3}, 3)$, and O for $(\mathbf{Opt}, 3)$. U_j^A will denote A's utility and U_j^l

[23] We described first the protocol for one encounter in detail since it is the basic component of the multi-encounter case.

[24] This is the situation when there are only two agents in the environment or when the type of one agent does not depend on the types of the others. That is, we do not consider the cases where by learning that an agent is of type h, for example, it can also conclude that others are less likely to be of the same type.

Short Notation	*Utility*	*Explanation*
U_l^l	$U^{W_l}(\hat{s}^{W_l,3}, 3)$	W_l's utility for agreement $\hat{s}^{W_l,3}$ in period 3
U_h^l	$U^{W_l}(\hat{s}^{W_h,3}, 3)$	W_l's utility for agreement $\hat{s}^{W_h,3}$ in period 3
U_O^l	$U^{W_l}(\mathbf{Opt}, 3)$	W_l's utility for opting out in period 3
U_l^A	$U^{A_j}(\hat{s}^{W_l,3}, 3)$	A_j's utility for agreement $\hat{s}^{W_l,3}$ in period 3
U_h^A	$U^{A_j}(\hat{s}^{W_h,3}, 3)$	A_j's utility for agreement $\hat{s}^{W_h,3}$ in period 3
U_O^A	$U^{A_j}(\mathbf{Opt}, 3)$	A_j's utility if W opts out in period 3

Figure 1: Short notation. Note that $U_O^A < U_h^A < U_l^A$ and $U_O^l < U_l^l < U_h^l$.

denotes W_l's utility. These notations are summarized in Figure 1.

Given the new definitions of histories, strategies and systems of beliefs, sequential equilibrium's three conditions impose restrictions on each time period in all sequences of histories. In particular, in any time period of every possible history in a sequence of histories, agent i will try to maximize its expected utility, with regard to the strategy of its opponent (which is composed of a sequence of strategies for each encounter) and its beliefs about the probabilities of its opponents' type according to the given history. Furthermore, if there is a probability attached as to whether a future encounter will happen, it will use this probability to compute its expected utility.

4 Two Agents Involved in Two Encounters

Suppose that there is some probability that the agents will meet again and negotiate in similar situations in the future. In particular, we assume that the agent that plays the role of A in the current negotiation, will play the same role in the future. This may happen in situations where there are agents that use the resource more often than the others. In some cases, A_j may take action in order to find out what W's type is. On the other hand, W may want to influence A_j's beliefs in order to benefit in future encounters and may be willing to lose now in order to increase its overall expected utility from all encounters. We demonstrate the problem using the following example.

Example 1
Suppose there are several automated bankers, and each has its own communication system to communicate with its customers and with other branches of the bank. In addition there is a public communication line available for payment. The automated bankers usually use their own communications systems; however, in the

case they are experiencing an overload, or in case their communications system is down, they use the public domain system. There are several possible standard contracts for using the public systems which are known to all agents; however, the specific contract of each bank is not known to the others.

Suppose that the communication system of BankA is down for several days and therefore it uses the public system. At the same time, while the private system of BankB is working, it is experiencing an exceptionally high workload. Therefore, BankB would like to use the public system in the next M time periods.

A negotiation ensues between the two agents over dividing the use of the public system, during which time BankA has sole access to the public system and continues to serve its customers during the negotiation, and BankB cannot serve any of its extra customers. The negotiation requires resources from both agents; however, BankA's gain from usage of the resource is higher than its loses from the negotiation itself. If an agreement on sharing the public system is not reached, BankB can threaten to opt out and disconnect BankA from the public system. If BankA is disconnected, each has some probability gaining access to the public system after some time.[25] *During the delay due to "opting out", BankA will lose the connection with all its customers; if it is able later to gain access, it will need to pay again for the connection. BankB can still use its own communication line to serve a portion of its customers.*

They both know that on the following day, the communication system of BankA will still be down and there is a probability β that BankB's system will again be overloaded (this probability is based on their expectation from the behavior of the stock market, etc.).

In this case, BankA which is using the communications system plays the role of A, while BankB plays the role of W. Since, BankA continues to serve its customers during the negotiation and its profits from using the public system is higher than its loss from the negotiation, it gains during the negotiation. Therefore, it doesn't care if the negotiation continues indefinitely. In addition, it will prefer the same agreement later rather than sooner and opting out later rather than sooner; thus, A0, A2 and A3 hold for BankA.

Since BankB can't use the resource during the negotiation, but the negotiation is costly, it loses over time. Therefore, it would prefer any action over disagreement and it will prefer any agreement sooner rather than later and opting out sooner rather than later; thus A0, A2 and A3 hold for BankB.

Both agents would like to use the public system in the relevant time period as much as possible, which justifies A1.

Suppose that both agents believe that with probability $0 \leq \beta \leq 1$ they will meet again in a similar situation in the future. We also assume that any given

[25] Similar situations occur in the usage of public phones by humans.

case	Conditions	Subconditions	Results
1	$U^{A_j}(\hat{s}^{W_h,3},3)-(\phi^j U_l^A+$ $(1-\phi^j)U_O^A)>$ $\phi^j\beta(U_l^A-U_h^A)$		$(\hat{s}^{W_h,3},3)$ in both encounters
2a	$U_h^A-(\phi^j U_l^A+$ $(1-\phi^j)U_O^A)<$ $\phi^j\beta(U_l^A-U_h^A)$	$U_l^l-U_O^l<$ $\beta(U_h^l-U_l^l)$	$(\hat{s}^{W_h,3},3)$ in both encounters
2b		$U_l^l-U_O^l>$ $\beta(U_h^l-U_l^l)$	If W_l then $(\hat{s}^{W_l,3},3)$ in both encounters. If W_h $(\mathbf{Opt},3)$ in the first encounter $(\hat{s}^{W_h,3},3)$ in the second

Table 1: Results when the expectation for A from $\hat{s}^{W_h,t}$ is higher than from $\hat{s}^{W_l,t}$ (Section 4.1.)

type of agent has the same beliefs concerning its opponent. We will denote by $\phi^j, j \in \mathbf{Type}$ A_j's beliefs that W is of type l. For example, all the agents A_h believe that their opponents are of type l with probability ϕ^h, and type h with probability $1-\phi^h$.[26]

In the next sections we will characterize the situations of negotiation by two agents in two encounters using different conditions. Note, that these conditions compliment one another, and provide us with a large range of situations. The conditions consist of inequalities on the utility functions of the agents. When the inequality is with respect to agent A the letter A will appear in the condition's title. For example, BA1 denotes an inequality involving agent's A's utility function. If an inequality is denoted by 1, its reverse will be denoted by 2. For example, the reverse of inequality BA1 is denoted by BA2. As we mentioned above, these two inequalities cover all the possibilities of A's utility functions, beside the one that yields equality.

4.1 A's expectation from $\hat{s}^{W_h,t}$ is higher than from $\hat{s}^{W_l,t}$

We first consider the case where A_j's expected utility (regardless of its type) in a single encounter from offering $\hat{s}^{W_h,t}$ is greater than offering $\hat{s}^{W_l,t}$ (when it is A's turn to make an offer). As was discussed in Section 3.4, if offered $\hat{s}^{W_h,t}$, both types of W will accept the offer, but if offered $\hat{s}^{W_l,t}$, W_h will opt out. Therefore, in the current section we assume that the following holds:

BA1 $\quad \forall j \in \mathbf{Type}, \phi^j U^{A_j}(\hat{s}^{W_l,t},t)+(1-\phi^j)U^{A_j}(\mathbf{O},t)<U^{A_j}(\hat{s}^{W_h,t},t).$

[26]This assumption is reasonable if there is a finite number of beliefs about other agents. Then, we can divide any type into subtypes according to its beliefs about its opponent.

In particular, the utility for A_j from offering $\hat{s}^{W_l,3}$ is U_l^A if W is of type l (which A_j believes with probability ϕ^j) and U_O^A if W is of type h (which A_j believes with probability $1-\phi^j$). Thus, the expected utility for A_j from offering $\hat{s}^{W_l,3}$ is smaller than the utility of offering $\hat{s}^{W_h,3}$ which will be accepted by W regardless of its type and will provide A_j with U_h^A with certainty in the second iteration of the negotiation.

The main question here is, if another encounter is possible with probability β, whether it is worthwhile for A_j to offer $\hat{s}^{W_l,3}$ in the first encounter. In such cases, if its opponent's type is l, A_j may find out about it and use its findings in the second encounter. In that case, A_j should compare between its expected utility from offering $\hat{s}^{W_h,3}$ in both encounters, i.e., $U_h^A + \beta U_h^A$, and offering $\hat{s}^{W_l,3}$ in the first encounter. After that A_j can then decide according to the results whether to offer $\hat{s}^{W_l,3}$ again or $\hat{s}^{W_h,3}$, i.e., $\phi^j[U_l^A + \beta U_l^A] + (1-\phi^j)[U_O^A + \beta U_h^A]$.

In the following theorem we consider the situation where the possible loss for A_j from offering $\hat{s}^{W_l,3}$ rather than $\hat{s}^{W_h,3}$ in the first encounter is greater than the possible gain for A_j from finding out that W is of type l and then reaching the agreement $\hat{s}^{W_l,3}$. Formally,

BA1.1 $\quad U_h^A - [\phi^j U_l^A + (1-\phi^j)U_O^A] > \phi^j \beta [U_l^A - U_h^A]$.

In the next theorem we show that if BA1.1 holds both encounters will always end with an agreement.

Theorem 6 (A_j does not gain enough from information)
If the model satisfies assumptions $A0-A6$, BA1 and BA1.1 and the agents use sequential equilibrium strategies then A_j should offer $\hat{s}^{W_h,3}$ in the second iteration of both encounters. This offer will be accepted by both types of W (this is **Case 1** *of Table 1).*

Proof
If A_j offers $\hat{s}^{W_h,3}$ in the first encounter, W will accept the offer regardless of its type. On the other hand, if A_j offers $\hat{s}^{W_l,3}$, it is clear from the discussion of Section 3.4 that W_h will opt out and W_l will consider accepting the agreement. If W_l accepts the agreement then A_j will realize that W's type is l and can use it in the second encounter (if it occurs). Thus, in order to conclude that A_j will offer $\hat{s}^{W_h,3}$ we need to show that: $U_h^A + \beta U_h^A \geq \phi^j[U_l^A + \beta U_l^A] + (1-\phi^j)[U_O^A + \beta U_h^A]$, but this is clear from the assumption BA1.1. $\square$

The equilibrium in the above theorem is a pooling equilibrium. Both types of W will take the same actions, and agent A will be able to obtain additional information on W in these encounters. We now consider situations in which the inequality of BA1.1 is reversed[27] and make additional assumptions about W_l's utility function.

BA1.2 $\quad U_h^A - [\phi^j U_l^A + (1-\phi^j) U_O^A] < \phi^j \beta [U_l^A - U_h^A]$

The next inequality states that the possible loss for W_l from opting out in the first encounter rather than accepting $\hat{s}^{W_l,3}$ is less than the possible gain for W_l in the second encounter given β, from offering U_h^l rather than U_l^l.

BW1 $\quad U_l^l - U_O^l < \beta[U_h^l - U_l^l]$.

If BA1.2 holds there may be situations where it is worthwhile for A_j to offer $\hat{s}^{W_l,3}$ (**Case 2** of Table 1), depending on W_l's behavior. If BW1 holds, it is worthwhile for W_l to pretend to be W_h by opting out when offered $\hat{s}^{W_l,3}$ in the first encounter to conceal its type. This will ensure that W will be offered $\hat{s}^{W_h,3}$ in the second encounter. This is stated in the following theorem.

Theorem 7 (A_j may benefit from information, but W_l conceals its type)
*If the model satisfies assumptions A0 – A6, BA1, BA1.2 and BW1 and the agents use sequential equilibrium strategies, then A_j will offer $\hat{s}^{W_h,3}$ in both encounters. (***Case 2a** *of Table 1.)*

Proof
If W_l is offered $\hat{s}^{W_l,3}$ in the first encounter, it should consider whether to accept the offer and thus reveal its type (because W_h will never accept this offer), or opt out and receive $\hat{s}^{W_h,3}$ in the next encounter too (if it occurs). That is, if $(1+\beta)U_l^l < U_O^l + \beta U_h^l$ then W_l should opt out in the first encounter if it receives $\hat{s}^{W_l,3}$. However, this could be concluded from BW1.

If W_l will opt out in the first encounter if offered $\hat{s}^{W_l,3}$, it is better for A_j to offer $\hat{s}^{W_h,3}$ in the first encounter (since A_j cannot learn anything by offering $\hat{s}^{W_l,3}$), so that both encounters will end with an agreement $\hat{s}^{W_h,3}$ in the second iteration. □

Next, we consider situations where the reverse of inequality BW1 holds.[28]

BW2 $\quad U_l^l - U_O^l > \beta[U_h^l - U_l^l]$.

If BW2 holds, it is worthwhile for W_l to accept $\hat{s}^{W_l,3}$ in the first encounter (**Case 2b** of Table 1). In this situation, it is worthwhile for A_j to offer $\hat{s}^{W_l,3}$ in the first encounter and to find out what W's type is.

[27] Note, that inequalities BA1.1 and BA1.2 below cover all possibilities of A's utility functions, beside the one that yields equality.

[28] Note, that inequalities BW1 and BW2 cover all possibilities of W_l's utility functions, beside the one that yields equality.

Theorem 8 (A_j may benefit from information, and W_l reveals its type)
If the model satisfies assumptions $A0 - A6$, BA1, BA1.2 and BW2, and the agents use sequential equilibrium strategies, then A_j will offer $\hat{s}^{W_l,3}$ in the first encounter and decide on its offer in the next encounter according to W's behavior in the first one. If W opts out, A_j will then offer $\hat{s}^{W_h,3}$ in the second encounter and if W accepts the offer in the first encounter, A_j will then offer it $\hat{s}^{W_l,3}$ again in the second encounter.

Proof
If $U_l^l - U_O^l \geq \beta(U_h^l - U_l^l)$ then it isn't worthwhile for W_l to opt out in the first encounter even if it receives an offer of $\hat{s}^{W_l,3}$. In such situations it is worthwhile for A_j to try to find out W's type by offering $\hat{s}^{W_l,3}$. This offer will be accepted by W_l, but W_h will opt out. $\square$

The equilibrium in the above theorem is a separating equilibrium. At the end of the first encounter A will find out what W's type is.

We demonstrate the different cases of this section in the following example.

Example 2
We return to the example of the communication systems. Suppose that BankA does not know the exact details of the contract BankB has for using the public communication system.[29] *We also assume that each negotiation iteration takes one minute. There are two possibilities for BankB's contracts: high (h) and low (l).*

If BankB holds contracts of type h, then the BankB loses \$300 per minute during the negotiation period and gains \$100 per minute when sharing the communication system with BankA. If BankB opts out the overall gain for BankB will be \$450, but it will also lose \$100 per minute during the negotiation.

If BankB holds contracts of type l, then BankB loses \$200 per minute during the negotiation period and gains \$100 per minute when sharing the public communication system with BankA. If BankB opts out the overall gain for BankB will be \$250, but it will also lose \$100 per minute during the negotiation as in type h.

BankA gains \$600 per minute during the negotiation and gains \$100 per minute when it shares the public communication system with the BankB. If BankB opts out of the negotiation BankA obtains \$300 per minute for the time of the negotiation.

The utility functions of the agents are presented in Table 2.

Suppose BankB would like to gain access to the public communication system for the next 10 time periods, i.e., $M = 10$ and the BankA believes that BankB is

[29] For simplicity, we assume that BankB's contract is known to BankA. However, it is easy to extend the example to the case where there are two types of contracts for BankA too.

Type h	*Type l*
$U^{bb_h}((s,t)) = s_u - 3\lfloor\frac{t}{2}\rfloor$	$U^{bb_l}((s,t)) = s_u - 2\lfloor\frac{t}{2}\rfloor$
$U^{bb_h}((\mathbf{Opt}_{bb},t)) = 4.5 - \lfloor\frac{t}{2}\rfloor$	$U^{bb_l}((\mathbf{Opt}_{bb},t)) = 2.5 - \lfloor\frac{t}{2}\rfloor$
$U^{bb_h}((\mathbf{Opt}_{ba},t)) = -\lfloor\frac{t}{2}\rfloor$	$U^{bb_l}((\mathbf{Opt}_{ba},t)) = -\lfloor\frac{t}{2}\rfloor$
$\hat{s}^{bb_h,t} = (5 - 2\lfloor\frac{t}{2}\rfloor, 5 + 2\lfloor\frac{t}{2}\rfloor)$	$\hat{s}^{bb_l,t} = (7 - \lfloor\frac{t}{2}\rfloor, 3 + \lfloor\frac{t}{2}\rfloor)$

$U^{ba}((s,t)) = s_{ba} + 6\lfloor\frac{t}{2}\rfloor$
$U^{ba}((\mathbf{Opt}_{bb},t)) = 3\lfloor\frac{t}{2}\rfloor$
$U^{ba}((\mathbf{Opt}_{ba},\lfloor\frac{t}{2}\rfloor)) = -100$

Table 2: Utility functions of the BankA and BankB in Example 2. bb_l and bb_h denote the BankB of type l and h respectively. ba denotes BankA.

of type l with probability $\frac{7}{12}$, i.e., $\phi = \frac{7}{12}$. BankA, of course, plays the role of A and BankB plays the role of W. This is a situation where A's expectation from $\hat{s}^{W_h,t}$ is higher than from $\hat{s}^{W_l,t}$ in a single encounter. In particular, $\hat{s}^{W_l,3} = (6,4)$, $\hat{s}^{W_h,3} = (3,7)$, $U_l^A = 12$, $U_h^A = 9$ and $U_O^A = 3$, $U_h^l = 5, U_l^l = 2$, and $U_O^l = 1.5$.

If the probability that BankB will need the public communication system again is, for example, $\frac{2}{7}$ (i.e., $\beta = \frac{2}{7}$) then it is easy to compute that we are in **Case 1** *of Table 1. In that case there will be an agreement in both encounters.*

However, if the probability that BankB will need the public communication system again is $\frac{4}{7}$ (i.e., $\beta = \frac{4}{7}$) then we are in **Case 2** *of Table 1. Given the utility function of BankB of type l, we are in* **Case 2a** *of Table 1 and therefore, this case will also end with an agreement in both encounters.*

To summarize, the cases that we consider cover all possibilities of utility functions when A's expectation from $\hat{s}^{W_h,t}$ is higher than from $\hat{s}^{W_l,t}$, beside situations of equality in the conditions. In all the cases we considered in this section, the negotiation will end in the second iteration in both encounters. Also, the second encounter will always end with an agreement. However, some of the first encounters will end with opting out (**Case 2b**). This seems to be a rare situation, since this may happen only if the probability of the second encounter (i.e., β) is very low and W_l's utilities from $\hat{s}^{W_l,3}$ and $\hat{s}^{W_h,3}$ are very close in value, considering that $U_l^l - U_O^l$ is very close to zero.

4.2 A_j's expectation from $\hat{s}^{W_h,t}$ is lower than from $\hat{s}^{W_l,t}$

In this section we consider the complementary case to Section 4.1 where, in a single encounter, the expected utility for A_j, regardless of its type, from offering $\hat{s}^{W_l,t}$ is greater than from offering $\hat{s}^{W_h,t}$, i.e.,

BA2 $\quad \phi^j U^{A_j}(\hat{s}^{W_l,t},t) + (1-\phi^j)U^{A_j}(\mathbf{O},t) > U^{A_j}(\hat{s}^{W_h,t},t)$

In this situation, if there is a single encounter, A_j will offer $\hat{s}^{W_l,t}$ and will always have the opportunity to get information on W's type. Therefore, in such situations, during the first negotiation encounter, W_l wants to convince A_j that it is of type h. If it succeeds, in the next encounter A_j will treat W as W_h and not as W_l as if its beliefs are not changed.[30] The only way that W_l may convince A_j that its type is h, is by opting out if it gets an offer less than $\hat{s}^{W_h,t}$. As we explained above, if W_h is offered less than $\hat{s}^{W_h,t}$ then it opts out since opting out is better for W_h than an offer that is less than $\hat{s}^{W_h,t}$. If W_h rejects the offer (chooses **No**) it will not reach a better agreement in the future. Therefore, if W_l wants to convince A_j that it is of type h it should opt out too. However, it is not always rational for W_l to pretend being W_h and to opt out. It depends on condition BW2 that is described in the previous section in which the difference between W_l's utility from $\hat{s}^{W_l,3}$ and $(\mathbf{Opt}, 3)$ is greater than the difference between $\hat{s}^{W_l,3}$ and $\hat{s}^{W_h,3}$ multiplied by β.

In particular, if BW2 holds, it is not rational for W_l to pretend being W_h as stated in the following theorem.

Theorem 9 (W_l accepts $\hat{s}^{W_l,3}$ and reveals its type)
If the model satisfies assumptions $A0-A6, BA2, BW2$, and the agents use sequential equilibrium strategies, if A_j will offer $\hat{s}^{W_l,3}$ in the second iteration of the first encounter, then W_h will opt out and will be offered $\hat{s}^{W_h,3}$ in the second encounter. W_l will accept $\hat{s}^{W_l,3}$ and will be offered the same in the second encounter.

Proof
If A_j knows that W is of type l it will offer $\hat{s}^{W_l,t}$ in any iteration of the negotiation. On the other hand, since W_h will not accept an offer of $\hat{s}^{W_l,t}$ if W accepts such an offer, A_j can conclude that it is of type l. Therefore, W_l should compare between accepting $\hat{s}^{W_l,3}$ in both encounters, or opting out in the first encounter and accepting $\hat{s}^{W_h,3}$ in the next one (if it occurs). Thus, if $U_l^l + \beta U_l^l \geq U_O^l + \beta U_h^l$, W_l should accept $\hat{s}^{W_l,3}$. But this inequality is clear from BW2 of the theorem. Now, since in this section we assume that A_j prefers to obtain information (i.e., BA2 holds), the theorem is clear. $\square$

As in Theorem 8 the equilibrium in the above theorem is separating equilibrium and A_j will find out what W's type is in the second encounter.

If BW1 holds rather than BW2, then the situation is more complicated. In these cases, W_l's expected utility from opting out in the second iteration of the first encounter, and accepting an offer in the next encounter as W_h, is greater than its expected utility from accepting an offer as W_l in both encounters. Even

[30] Note that in the previous section, W_l just wanted to maintain the current situation. In this case, W_l tries to convince A_j that its type is h.

though it is worthwhile for W_l to opt out in the first encounter, if it is offered $\hat{s}^{W_l,3}$ in the second iteration and if as a result A_j will change its beliefs and offer W $\hat{s}^{W_h,3}$ in the second encounter. However, if W_l will always behave as W_h, then A_j will not change its beliefs if it observes behavior typical to W_h, since it knows that both W_l and W_h behave similarly[31], i.e., opting out when offered $\hat{s}^{W_l,3}$. Thus, there is no sequential equilibrium with pure strategies, and the agents should use mixed strategies and the equilibrium is hybrid equilibrium.

As we mentioned in Section 3.2, when the agents choose to randomize between several pure strategies, the expected utility from all of these strategies should be the same, otherwise they will not agree to randomize but rather prefer one pure strategy over the other. In our case, when W_l is offered $\hat{s}^{W_l,3}$ it should randomize between accepting the offer and opting out. Thus, its expected utility in both cases should be the same. If A_j observes opting out, it should randomize in the second encounter between offering $\hat{s}^{W_l,3}$ again or offering $\hat{s}^{W_h,3}$.

We denote by p_W the probability that W_l will opt out if offered $\hat{s}^{W_l,3}$ in the first encounter, and by p_A the probability that A_j will offer $\hat{s}^{W_h,3}$ in the second encounter if W opts out of the first one. The probability in which W_l and A_j should randomize their strategies in a sequential equilibrium is stated in the next lemma.

Lemma 10 (probabilities of mixed strategies of W_l and A_j)
If the model satisfies assumptions $A0-A6, BA2, BW1$ and the agents use sequential equilibrium strategies (with mixed strategies) then

$$p_A = \frac{U_l^l - U_O^l}{\beta(U_h^l - U_l^l)} \tag{1}$$

and

$$p_W = \frac{(1-\phi^j)(U_h^A - U_O^A)}{\phi^j(U_l^A - U_h^A)} \tag{2}$$

Proof
Suppose that if W_l receives an offer $\hat{s}^{W_l,3}$ in the first encounter, it will opt out with probability p_W. Since W_h always opts out in such situations, when A_j observes opting out after proposing $\hat{s}^{W_l,3}$, it will update its beliefs about W's type and using Bayes' rule it will conclude that W's type is l with probability $\frac{\phi^j p_W}{1-\phi^j+\phi^j p_W}$. In the next encounter (if it occurs), A_j will choose randomly between offering $\hat{s}^{W_l,3}$ and $\hat{s}^{W_h,3}$ if the expected utilities from both offers are the same. If it

[31] Note that there is a paradox. If A_j is not influenced by both agents behaving as W_h, it is not rational for W_l to behave as W_h. Therefore, if A_j observes W_h's behavior, it may conclude that it is W_h. However, if A_j's beliefs are affected by W's behavior, it is worthwhile for W_l also to pretend to be W_h; a paradox.

offers $\hat{s}^{W_h,3}$, then W will accept the offer regardless of its type and A_j's expected utility will be U_h^A. If it offers $\hat{s}^{W_l,3}$ in the second encounter, W_l will accept the offer and W_h will opt out. Using A_j updated belief, its expected utility is: $\frac{\phi^j p_W}{1-\phi^j+\phi^j p_W} U_l^A + (1 - \frac{\phi^j p_W}{1-\phi^j+\phi^j p_W}) U_O^A$. If A_j expected utilities from both $\hat{s}^{W_h,3}$ and $\hat{s}^{W_l,3}$ to be the same we can conclude that $\frac{\phi^j p_W}{1-\phi^j+\phi^j p_W} U_l^A + (1 - \frac{\phi^j p_W}{1-\phi^j+\phi^j p_W}) U_O^A = U_h^A$ and thus $p_W = \frac{(1-\phi^j)(U_h^A - U_O^A)}{\phi^j (U_l^A - U_h^A)}$.

If W_l receives the offer $\hat{s}^{W_l,3}$ in the first encounter, it will choose randomly between opting out and accepting the offer only if its expected utility from both are the same. If it accepts the offer and reveals its type, it will be offered $\hat{s}^{W_l,3}$ also in the next encounter, thus W_l's expected utility in this case is $U_l^l + \beta U_l^l$. If it opts out with probability p_W, then in the second encounter A_j will offer $\hat{s}^{W_h,3}$ with probability p_A and will offer $\hat{s}^{W_l,3}$ with probability $1 - p_A$. Thus, W_l's expected utility in this case is $U_O^l + \beta(p_A U_h^l + (1 - p_A) U_l^l)$. We require that $U_O^l + \beta(p_A U_h^l + (1 - p_A) U_l^l) = U_l^l + \beta U_l^l$ and we conclude that $p_A = \frac{U_l^l - U_O^l}{\beta(U_h^l - U_l^l)}$.

It is still left to be shown that $0 \leq p_A \leq 1$ and that $0 \leq p_W \leq 1$. From our assumptions, $U_l^l > U_O^l$ and $U_h^l > U_l^l$ it is clear that $p_A > 0$ and from BW1 it is clear that $p_A < 1$. Similarly, since $U_h^A > U_O^A$ and $U_l^A > U_h^A$ $p_W > 0$ and by BA2 it is clear that $p_W < 1$. □

Finally, we must verify that under the above mixed strategies, it is still worthwhile for A_j to offer $\hat{s}^{W_l,3}$ in the first encounter, where W_h will opt out and W_l will choose randomly between also opting out or accepting. This is considered in the following theorem.

Theorem 11 (mixed strategies and pure strategies)
If the model satisfies assumptions $A0 - A6, BA2, BW1$ and the agents use sequential equilibrium strategies (with mixed strategies) then if

> **BA23** :
> $$\phi^j(1 - p_W)(U_l^A + \beta U_l^A) +$$
> $$(1 - \phi^j + \phi^j p_W)[U_O^A + \beta p_A U_h^A + \beta(1 - p_A)(\phi^j U_l^A + (1 - \phi^j) U_O^A)] >$$
> $$U_h^A + \beta(\phi^j U_l^A + (1 - \phi^j) U_O^A)$$

then

First encounter: *A_j will offer $\hat{s}^{W_l,3}$ in the first encounter; W_h will always opt out and W_l will opt out with probability $p_W = \frac{(1-\phi^j)(U_h^A - U_O^A)}{\phi^j(U_l^A - U_h^A)}$ and with probability $1 - p_W$ will accept the offer. If W accepts the offer, agent A will believe with probability one, that W's type is l. If W opts out, A will believe with probability $\frac{U_h^A - U_O^A}{U_l^A - U_O^A}$ that W's type is l.*

Second encounter: *If A believes that W's type is l with probability 1, then it will offer $\hat{s}^{W_l,3}$ which be accepted by W.*[32]

Otherwise, A_j will offer $\hat{s}^{W_l,3}$ with probability $p_A = \frac{U_l^l - U_O^l}{\beta(U_h^l - U_l^l)}$ and with probability $1 - p_A$ it will offer $\hat{s}^{W_h,3}$. W_l will accept the offer, but W_h will opt out.

If inequality BA23 does not hold, A_j will offer $\hat{s}^{W_h,3}$ in the first encounter and $\hat{s}^{W_l,3}$ in the second one.

Proof

Most of the proof is clear from Lemma 10 and the discussion in Section 3.4. It is left to be shown that if inequality BA23 holds, then A_j will offer $\hat{s}^{W_l,3}$ in the first encounter. According to Lemma 10, if A_j offers $\hat{s}^{W_l,3}$ in the first encounter, then it believes that with probability $\phi^j(1 - p_W)$ its offer will be accepted, W_l reveals its type and its overall expected utility in this case is $(U_l^A + \beta U_l^A)$.

A_j also believes that with probability $1 - \phi^j + \phi^j p_W$ W will opt out (either because it is W_h or because it is W_l that opts out with probability p_W). In this case its utility in the first encounter will be U_O^A and in the second encounter A_j will offer $\hat{s}^{W_h,3}$ with probability p_A and with probability $1 - p_A$ will offer $\hat{s}^{W_l,3}$. If it offers $\hat{s}^{W_h,3}$, its offer will be accepted by both agents; however, W_h will opt out if offered $\hat{s}^{W_l,3}$. To summarize, A_j's expected utility from offering $\hat{s}^{W_l,3}$ in the first encounter is:

$$\phi^j(1-p_W)(U_l^A+\beta U_l^A)+(1-\phi^j+\phi^j p_W)[U_O^A+\beta p_A U_h^A+\beta(1-p_A)(\phi^j U_l^A+(1-\phi^j)U_O^A)]$$

If A_j offers $\hat{s}^{W_h,3}$ it will be accepted by W regardless of its type, and A_j's beliefs will not be changed. According to BA2 in the second encounter, A_j will offer $\hat{s}^{W_l,3}$, thus, A_j's expected outcome in this case is $U_h^A + \beta(\phi^j U_l^A + (1 - \phi^j)U_O^A)$.

Therefore, A_j will offer $\hat{s}^{W_l,3}$ in the first encounter if the following holds:

$$\begin{aligned}
&(1 - \phi^j p_W)(U_l^A + \beta U_l^A)+\\
&(1 - \phi^j + \phi^j p_W)[U_O^A + \beta p_A U_h^A + \beta(1 - p_A)(\phi^j U_l^A + (1 - \phi^j)U_O^A)] >\\
&U_h^A + \beta(\phi^j U_l^A + (1 - \phi^j)U_O^A.
\end{aligned}$$

By BA23 this inequality holds. □

It is useful to characterize situations where the condition BA2.3 of the above theorem holds. Especially since p_A includes W_l's utility, it is useful to know whether A_j's decision depends on W_l's utility or not. We found out that A_j's decision depends on its own utilities, A_j's original belief that W's type is l (ϕ^j) and on the probability that the agents will meet again (β).

[32] Note, that A's belief is correct here in this case; W's type is l.

Lemma 12
If the model satisfies assumptions $A0-A6, BA2, BW1$ and the agents use sequential equilibrium strategies (with mixed strategies) and

$$U_l^A - U_h^A + (1-\phi^j)U_O^A + \beta(1-\phi^j)U_l^A > (1-\phi^j)\frac{(U_h^A - U_O^A)(U_l^A - U_O^A)}{U_l^A - U_h^A} \quad (3)$$

then A_j will offer $\hat{s}^{W_l,3}$ in the first encounter.

Proof
After substituting p_W and p_A according to their definitions in inequality 3 we obtain the following:

$$U_h^A + \beta U_h^A - (1 - \tfrac{(1-\phi^j)(U_h^A-U_O^A)}{U_l^A-U_h^A})(U_l^A + \beta U_l^A) \quad (4)$$
$$-(1-\phi^j + \tfrac{(1-\phi^j)(U_h^A-U_O^A)}{U_l^A-U_h^A})(U_O^A + \tfrac{\beta U_h^A(U_h^l-U_O^l)}{\beta(U_h^l-U_l^l)}) + \beta(1 - \tfrac{U_l^l-U_O^l}{\beta(U_h^l-U_l^l)}(\phi^j U_l^A + (1-\phi^j)U_O^A)$$
$$< 0.$$

After some manipulations of the above, one can conclude that inequality 3 holds.
□

In the following examples we demonstrate the situations described in this section.

Example 3
We return to the example of the communication systems. Suppose that the utility functions are exactly as in Table 2 of Example 2, but that BankA believes with probability $\frac{3}{4}$ that the BankB's type is l. We have $U_l^A = 12, U_h^A = 9, U_O^A = 3, U_h^l = 5, U_l^l = 2$, and $U_O^l = 1.5$ In this situation we are in the case where A's expectation from $\hat{s}^{W_h,t}$ is lower than from $\hat{s}^{W_l,t}$ in a single encounter, i.e., BA2 holds.

Suppose $\beta = \frac{1}{10}$. In this situation BW2 holds and by Theorem 9 BankA (A) will offer $\hat{s}^{W_l,3} = (6,4)$ in the first encounter. If BankB is of type l, it will accept the offer $\hat{s}^{W_l,3} = (6,4)$ and will get a similar offer in the second encounter (if it occurs). If BankB is of type h it will opt out in the first encounter, and will accept $(3,7)$ in the second encounter. In both cases, at the end of the first encounter BankA will know for sure the type of BankB.

Suppose, $\beta = \frac{1}{2}$. In this situation BW1 holds and the reverse of inequality BA23 is true. Therefore, by Theorem 11 BankA will offer $(3,7)$ in the first encounter and will offer $(6,4)$ in the second encounter.

Suppose, $\beta = \frac{1}{2}$ as before, but $U_l^A = 20$. In this situation BW1 still holds but inequality BA23 also holds. By Theorem 11 if BankB's type is l and it is offered $(6,4)$ in the first encounter it should choose randomly between accepting *the offer and opting out in the first encounter; with probability* $p_W = \frac{2}{5}$ *it will opt out and with probability* $\frac{3}{5}$ *it will accept the offer. BankA will offer* $(6,4)$ *in the first encounter. In the second encounter, with probability* $p_A = \frac{1}{3}$ *it will offer* $(3,7)$ *and with* $\frac{2}{3}$ *it will offer* W $(6,4)$.

In the previous section, where BA1 holds, we were able to find pure sequential equilibrium strategies, while in this section, where BA2 holds, in some situations the agents will need to use mixed strategies. The reason for this behavior is that when BA1 holds, W_l does not try to change A's belief in the first encounter, it just tries not to reveal its type. However, when BA2 holds, W_l tries to decrease A's probabilistic belief that its type is l. In several of these situations there is no sequential equilibrium of pure strategies.

5 Two Agents and More Than Two Encounters

In this section we consider situations in which the agents may meet more than twice.[33] For example, there is some probability that BankB's system will again be overloaded while the communication system of BankA is still down.

Suppose the agents believe that in addition to the current encounter there is some positive probability for m encounters and that the (independent) probability for each of these encounters is β_i, $i = 1, ..., m$, respectively. This assumption is valid in the case that the probability of the need for a resource in one time period does not depend on the probability of using the resource in another time period. For example, the probability that the communication system will be down on one day does not depend on whether it was down on the previous day. Similarly, the fact that there are excessive customer requests on one day does not help in predicting the situation on the next day. It mainly depends on the behavior of the market, the day of the week, etc. Thus, in our analysis, agents don't update their beliefs about future interactions.

As in the case of two encounters, we assume that agents play the same role in all encounters and again we will distinguish between the case where A_j's expectation from $\hat{s}^{W_h,t}$ is higher than from $\hat{s}^{W_l,t}$ in a single encounter, and the case where A_j's expectation from $\hat{s}^{W_h,t}$ is lower than from $\hat{s}^{W_l,t}$ in a single encounter.

5.1 *A_j's expectation from $\hat{s}^{W_h,t}$ is higher than from $\hat{s}^{W_l,t}$ (many encounters)*

We assume that BA1 holds. In this case, as in the case of two encounters, A_j may try to offer $\hat{s}^{W_l,t}$ in some of the encounters in order to find out what W's

[33] Since the case of two encounters is much simpler than the general case and the strategies are simpler, we decided to present both results.

type is. A_j's benefits from finding that W's type is l increases with the number of encounters. Therefore, A_j should compare between offering $\hat{s}^{W_h,t}$ in all encounters (an offer which will be accepted by both types of W), and offering $\hat{s}^{W_l,t}$ in the first encounter (which will cause W_h to opt out) and depending on the result, offer $\hat{s}^{W_l,t}$ or $\hat{s}^{W_h,t}$ in later encounters. We first consider the case where the expected loss from offering $\hat{s}^{W_l,3}$ rather than $\hat{s}^{W_h,3}$ in the first encounter, is greater than the expected gain from the information on W's type in possible future encounters. This is an extension of assumption BA1.1 for more than two encounters and it is formalized as follows:

BA1.1M $U_h^A - [(\phi^j U_l^A + (1-\phi^j)U_O^A)] > (\sum_{q=1}^m \beta_q)[(\phi^j U_l^A + (1-\phi^j)U_h^A) - U_h^A]$.

If BA1.1M holds, A_j should offer W $\hat{s}^{W_h,3}$ in both encounters as stated in the following theorem.

Theorem 13 (A_j does not gain enough from information)
If the model satisfies assumptions $A0-A6$, BA1, B1.1M and the agents are using sequential equilibrium strategies, then A_j should offer $\hat{s}^{W_h,3}$ in all the encounters. This offer will be accepted by W regardless of its type.

Proof
A_j should compare between the expected outcome from offering $\hat{s}^{W_l,3}$, in which case it would have an opportunity to find out what W's type is, and will have the following expected utility: $\phi^j U_l^A(1+\beta_1+\cdots+\beta_m)+(1-\phi^j)[U_O^A+U_h^A(\beta_1+\cdots+\beta_m)]$, and offering $\hat{s}^{W_h,3}$ in all encounters, in which case its expected utility will be $(1+\beta_1+\cdots+\beta_m)U_h^A$. Putting this together, the result will be that if the following inequality holds, A_j should offer $\hat{s}^{W_h,3}$ in all encounters:

$$(1+\sum_{i=1}^m \beta_i)U_h^A > \phi^j[U_l^A + (\sum_{i=1}^m \beta_i)U_l^A] + (1-\phi^j)[U_O^A + (\sum_{i=1}^m \beta_i)U_h^A] \qquad (5)$$

But, this is clear from assumption BA1.1M. □

We demonstrate this case in the following example.

Example 4
Suppose the situation is exactly as in the first case of Example 2, but in this case there are four possible encounters, in addition to the first one (i.e., $m=4$), and the probability of each of the encounters is $\beta = 0.1$. Note that the probability of each encounter is independent of the others; thus if BankA and BankB negotiate on the second day, it doesn't influence their probability of negotiating on the third day. In such a case BA1 and BA1.1M hold, and by Theorem 13, A_j will offer $\hat{s}^{W_h,3} = (3,7)$ in all the encounters, and the offer will be accepted by W.

We now consider the situation where the inequality of BA1.1M is reversed. In this case it is an extension of the situation of two encounters where BA1.2 holds.

BA1.2M $U_h^A - [(\phi^j U_l^A + (1-\phi^j) U_O^A)] < (\sum_{q=1}^m \beta_q)[(\phi^j U_l^A + (1-\phi^j) U_h^A) - U_h^A]$

If BA1.2M holds, A_j may consider offering W $\hat{s}^{W_l,3}$, but that depends on W_l's response to such an offer. As in the two encounter case, it may be that W_l, if offered $\hat{s}^{W_l,3}$, will prefer to select randomly between opting out and accepting the offer in order to prevent A_j from changing its beliefs. This will be rational for W_l if opting out and then receiving $\hat{s}^{W_h,3}$ in all the encounters that are left is better than receiving $\hat{s}^{W_l,3}$ in all encounters.

Otherwise, this situation is similar to the situation when BW2 holds and there were only two encounters. We denote the extension of BW2 by BW2M.

BW2M $(\sum_{q=1}^m \beta_q)(U_h^l - U_l^l) < U_l^l - U_O^l$.

Theorem 14 (A_j may benefit from information, and W_l reveals its type)
If the model satisfies assumptions $A0 - A6$, BA1, BA1.1M and BW2M and the agents use sequential equilibrium strategies, then in the first encounter A_j will offer $\hat{s}^{W_l,3}$. W_l will accept $\hat{s}^{W_l,3}$, W_h will opt out and A_j will update its beliefs accordingly. In the rest of the encounters if W is of type l, A_j will offer $\hat{s}^{W_l,3}$; otherwise it will offer $\hat{s}^{W_h,3}$.

Proof
If W_l's gain from accepting $\hat{s}^{W_h,3}$ instead of $\hat{s}^{W_l,3}$ in the rest of the encounters is smaller than W_l's loss from opting out instead of accepting $\hat{s}^{W_l,3}$ in the first encounter, it should accept $\hat{s}^{W_l,3}$ in the first encounter. But this is clear from assumption BW2M. □

Situations where the conditions of Theorem 14 hold are rare since $U_l^l - U_O^l$ is close to zero. They are true only when the probabilities of future encounters are low and when $U_h^l - U_l^l$ is relatively small. We present such a situation in the following example.

Example 5
Suppose the situation is as in the first case of Example 2 but W_l's utility from $\hat{s}^{W_h,3}$ is only three, i.e., $U_l^A = 12$, $U_h^A = 9$, $U_O^A = 3$, $U_h^l = 3$, $U_l^l = 2$, $U_O^l = 1.5$ and $\phi^j = \frac{7}{12}$. In addition, suppose there are three expected encounters (in addition to the first one) with probability $\beta = 0.15$. In this situation BA1,BA1.1M and BW2M hold. A_j will then offer $\hat{s}^{W_l,3}$ in the first encounter, W_h will opt out and W_l (which A_j believes with high probability is actually W's type) will accept the offer. The rest of the encounters will end either by $\hat{s}^{W_h,3}$ or $\hat{s}^{W_l,3}$ depending on W's real type.

The situation is more complicated if BW2M's inequality is reversed. This is an extension of the two encounters case where BW1 holds. Since $U^l_l - U^l_O$ is close to zero, it is a common situation.

BW1M $(\sum_{q=1}^{m} \beta_q)(U^l_h - U^l_l) > U^l_l - U^l_O$.

In such situations we will use backtracking techniques to identify sequential equilibrium. For this purpose we identify two types of encounters.

Definition 15 (maximal encounter in which it is still worthwhile for A_j to gather information)
Let n_A be the maximal encounter $0 \leq n_A \leq m$ which satisfies the following inequality:

$$U^A_h - [(\phi^j U^A_l + (1-\phi^j)U^A_O)] < (\sum_{q=n_A+1}^{m} \beta_q)[(\phi^j U^A_l + (1-\phi^j)U^A_h) - U^A_h] \qquad (6)$$

That is, n_A is the maximal encounter in which it is still worthwhile for A_j to try to gather information by offering $\hat{s}^{W_l,3}$ to W. It is clear that such an n_A exists (e.g., $n_A = 0$) since we consider the case where BA1.2M holds. It is also clear that $n_A \neq m$ since we consider the case in which A_j's expectation from $\hat{s}^{W_h,t}$ in a single encounter is higher than that from $\hat{s}^{W_l,t}$, i.e., BA1 holds.

We will consider A and W's behavior in encounter $n_A + 1$ and the encounters that came before it. At encounter $n_A + 1$ it is clear that if A_j's beliefs were not changed by it, so that A_j will offer $\hat{s}^{W_h,3}$ (since n_A is the maximal encounter in which it is still worthwhile for A to try to gather information regardless of its type) and W will accept it.

A_j's behavior before encounter n_A+1 depends on W_l's behavior which in turn depends on the maximal encounter so that it is still worthwhile for W_l to pretend being W_h.

Definition 16 (maximal encounter in which it is still worthwhile for W_l to pretend being W_h)
We denote by n_W the maximal encounter so that the following holds:

$$\left(\sum_{q=n_W+1}^{m} \beta_q\right)(U^l_h - U^l_l) > U^l_l - U^l_O \qquad (7)$$

That is, n_W is the maximal encounter in which it is still worthwhile for W_l to pretend being W_h by opting out if it receives an offer of $\hat{s}^{W_l,3}$. Note, that since in this case we have BW1M, then inequality 7 holds for $n_W = 0$.

The main factor that affects the sequential equilibrium strategies in this case is the relation between n_A and n_W. We will consider the case in which $n_A \leq n_W$. This is a common situation when $U_l^l - U_O^l$ is close to zero and in the presence of the other assumptions we make in this section. If $n_A \leq n_W$, then all encounters will end with an agreement.

Theorem 17 (W_l may wait longer than A_j)
If the model satisfies assumptions $A0 - A6, BA1, BA1.2M$ and $BW1M$, and the agents use sequential equilibrium strategies, then if $n_A \leq n_W$, agent A_j will offer $\hat{s}^{W_h,3}$ in all encounters, and it will be accepted by W regardless of its type.

Proof
We will show that for every encounter k, if it occurs, if A_j's beliefs were not changed before, and if it still believes with probability ϕ^j that W's type is l, then A_j should offer $\hat{s}^{W_h,3}$ and its beliefs will not change at the end of this encounter.
Base case: When $k > n_A$, this assumption is clear, because if A_j's beliefs were not changed, and since n_A is the maximal encounter in which it is still worthwhile for A_j to try and find out W's type by offering $\hat{s}^{W_l,3}$, if the encounter is greater than n_A, A_j should offer $\hat{s}^{W_h,3}$. This offer will be accepted by W, regardless of its type, and A_j's beliefs will not be changed.
When $k = n_A$, A_j should consider offering $\hat{s}^{W_l,3}$. However, since $n_W \geq n_A$ and for all encounters which are greater than n_A, it is clear that A_j will offer $\hat{s}^{W_h,3}$ if its beliefs were not changed. Therefore if W_l is offered $\hat{s}^{W_l,3}$, it should opt out and as a result, A_j's beliefs will not be changed and in the rest of the encounters it will offer $\hat{s}^{W_h,3}$. However, if this occurs, then A_j will not gain anything by offering $\hat{s}^{W_l,3}$ and it should actually offer $\hat{s}^{W_h,3}$.
Induction case ($k < n_A$): Suppose the assumption is true for all $k < k' \leq m$. In encounter k, if it occurs, A_j should consider offering $\hat{s}^{W_l,3}$. However, by the inductive assumption, if A_j's beliefs will not be changed then in the rest of its encounters A_j will offer $\hat{s}^{W_h,3}$. Since $k < n_W$, it is clear that it is worthwhile for W_l to opt out if offered $\hat{s}^{W_l,3}$. In this case, A_j will not benefit from offering $\hat{s}^{W_l,3}$ and it will therefore offer $\hat{s}^{W_h,3}$ which will be accepted by W, leaving A_j's beliefs the same as before.

We can conclude, that in all encounters, A_j will offer $\hat{s}^{W_h,3}$, which will be accepted by W regardless of its type. □

Note that the equilibrium in the above theorem is a pooling equilibrium. Both W_l and W_h will take the same actions. We demonstrate this case with the following example.

Example 6
Suppose the situation is similar to Example 2 but there are four possible encounters (in addition to the first one) each with probability $\beta = 0.2.$, i.e., $U_l^A = 12, U_h^A = 9, U_O^A = 3, U_h^l = 5, U_l^l = 2$, $U_O^l = 1.5$. and $\phi^j = \frac{7}{12}$. It is easy to see that BA1, BA1.2M and BW1M hold. Furthermore, if we denote the first encounter by 0, (thus the last one is denoted by 4), $n_A = 1$ and $n_W = 3$. By the above theorem, A_j will then offer $\hat{s}^{W_h,3} = (3,7)$ in all encounters.

5.2 A_j's expectation from $\hat{s}^{W_h,t}$ is lower than from $\hat{s}^{W_l,t}$ (many encounters)

In this section we assume that BA2 holds, and therefore A_j will always consider offering $\hat{s}^{W_l,3}$ and will gather information.

The simple case is when BW2M holds and it is *not* worthwhile for W_l to pretend to be W_h. The equilibrium in this case is a separating equilibrium and A_j finds out what is W's type. It is considered in the next theorem.

Theorem 18 (A_j may benefit from information and W_l reveals its type)
If the model satisfies assumptions $A0 - A6$, BA2 and BW2M, and the agents use sequential equilibrium strategies, then A_j will offer $\hat{s}^{W_l,3}$ in the first encounter which will be accepted by W_l and W_h will opt out. In future encounters, if it is W_h the agreement will be $\hat{s}^{W_h,3}$, and if it is W_l, the agreement will be $\hat{s}^{W_l,3}$.

Proof
Similar to the proof of Theorem 14. □

As we mentioned before, situations where BW2M holds are rare since $U_l^l - U_O^l$ is close to zero. We present such a situation in the next example.

Example 7
Suppose the agents' utility function is as in Example 5, i.e., $U_l^A = 12$, $U_h^A = 9$, $U_O^A = 3$, $U_h^l = 3$, $U_l^l = 2$, $U_O^l = 1.5$ but $\phi^j = \frac{3}{4}$. In addition, suppose that there are three expected encounters (in addition to the first one) with probability $\beta = 0.15$. In this case BA2 and BW2M hold and, according to the above theorem, if W is of type l then all encounters will end with $\hat{s}^{W_l,3}$. If W is of type h, it will opt out in the first encounter, but the rest of the encounters will end with $\hat{s}^{W_h,3}$.

If BW1M holds, then the situation is more complicated and the agents use mixed strategies. Usually, in the first encounter A_j will offer $\hat{s}^{W_l,3}$, W_h will opt out and W_l will choose randomly between opting out and accepting the offer. If $\hat{s}^{W_l,3}$ is accepted by W, then A_j will continue to offer $\hat{s}^{W_l,3}$ in the rest of the encounters. If W opts out, A_j's belief that W's type is h increases and in the next encounter it will choose randomly between offering $\hat{s}^{W_l,3}$ again and offering $\hat{s}^{W_h,3}$. W_l will choose randomly again between opting out and accepting and so on and so forth. Eventually, in the last encounter, if W_l is offered $\hat{s}^{W_l,3}$ it will accept it and W_h will opt out. Let us denote the first encounter by 0 and the rest of them by $1, \cdots, m$ as well as denoting by p_i $i = 0, \cdots, m-1$ the probability that W_l will opt out in encounter i if it receives $\hat{s}^{W_l,3}$. We denote by q_i $i = 1, \cdots, m$ the probability that A_j will offer $\hat{s}^{W_h,3}$ in encounter i if it chooses randomly between $\hat{s}^{W_l,3}$ and $\hat{s}^{W_h,3}$. Since, whenever an agent chooses randomly between two options, its expected outcome from both options should be the same. We can construct $2m$ equations specifying these equalities. Solving these equations, if possible, will provide us with the appropriate probabilities. However, these equations may turn out to be complicated because we should take into consideration whether a future encounter will occur and other possible scenarios. Both A_j's beliefs over time and the agents' expected utilities depend on the p_is and the q_is.

We prove a general lemma on A_j's beliefs concerning W's type in a given encounter.

Lemma 19 (A_j's belief after several encounters)
If the model satisfies assumptions $A0 - A6$, $BA2$ and $BW1M$, and the agents use sequential equilibrium strategies such that in encounter $0 \leq i < m$, if W_l is offered $\hat{s}^{W_l,3}$ and it opts out with probability p_i, then the following holds:

- *Suppose that the agents reach encounter $0 < y \leq m$ and before this encounter, encounters $0, i_1, \cdots, i_n$ occur in which W opts out, then if A_j's original belief that W's type is l was ϕ^j, and according to W_l's sequential equilibrium it opts out with probability p_{i_k} for $i_k = 0, \cdots, i_n$, then in the beginning of encounter y, A_j's belief that W's type is l is $\frac{\phi^j p_0 p_{i_1} \cdots p_{i_n}}{1-\phi^j+\phi^j p_0 p_{i_1} \cdots p_{i_n}}$.*
- *If A_j offered W $\hat{s}^{W_h,3}$ in a given encounter, then its belief in the next encounter will not change. Also, if a given encounter did not occur, A_j's belief does not change.*

Proof
The second conclusion is clear by our definitions. We prove the first item of the lemma by induction on i_n.
Base case (y=1): It is easy to see by Bayes' rule that after the first encounter if W opted out, A_j's belief is $\frac{\phi^j p_o}{1-\phi^j+\phi^j p_0}$.
Induction case: Suppose the assumption is true for encounter y and suppose W opted out in that encounter, then in the beginning of the following encounter,

A_j's beliefs will be according to Bayes's rule :

$$\frac{p_y \frac{\phi^j p_0 p_{i_1} \cdots p_{i_n}}{1-\phi^j+\phi^j p_0 p_{i_1} \cdots p_{i_n}}}{1 - \frac{\phi^j p_0 p_{i_1} \cdots p_{i_n}}{1-\phi^j+\phi^j p_0 p_{i_1} \cdots p_{i_n}} + p_y \frac{\phi^j p_0 p_{i_1} \cdots p_{i_n}}{1-\phi^j+\phi^j p_0 p_{i_1} \cdots p_{i_n}}}$$

Using simple algebraic manipulations we obtain the required probability. □

To demonstrate the kind of equations that should be considered, and to allow the reader to follow the reasoning, we will concentrate on a case restricted case which is specified by the following condition.

*A*7 **Three Encounters:** There are only three possible encounters and the second and third encounters have the same probability which is denoted by β.

We first consider whether it is beneficial for W_l to always threaten opting out in the first encounter if offered $\hat{s}^{W_l,3}$. In such a case, A_j's belief will not change, and in the second encounter the situation will be as in Theorem 11 of Section 4.2. However, it turns out that this threat is credible only if BW1 and the reverse of the inequality BA23 of Theorem 11 hold. In the other situations W_l prefers to choose randomly between accepting the offer of $\hat{s}^{W_l,3}$ and thereby revealing its type and opting out, and maintaining the situation.

Lemma 20 (W_l's behavior in the first encounter)
If the model satisfies assumptions $A0-A7$, BA2, BW1, BW1M and the agents use sequential equilibrium strategies then

- *If the* reverse *of inequality BA23 of Theorem 11 holds, then W_l will threaten to opt out whenever it is offered $\hat{s}^{W_l,3}$. Therefore, A_j will offer $\hat{s}^{W_h,3}$ in the first and second encounters and will offer $\hat{s}^{W_l,3}$ in the last encounter.*
- *If inequality BA23 of Theorem 11 holds, when offered $\hat{s}^{W_l,3}$, W_l will not threaten to opt out but may choose randomly between accepting the offer and opting out.*

Proof
If W_l accepts $\hat{s}^{W_l,3}$ in the first encounter and reveals its type, its expected utility is: $U_l^l + 2\beta U_l^l$.

If W_l decides to opt out in the first encounter it will receive U_O^l. The rest of its expected outcome depends on whether the second and third encounters will occur and how A_j will behave then.

We first consider the case where the reverse condition BA23 of Theorem 11 holds. In such a case if W_l opts out in the first encounter, then, if the second encounter occurs, A_j will offer $\hat{s}^{W_h,3}$. In the third encounter, A_j will always offer $\hat{s}^{W_l,3}$ which will be accepted by W_l. Thus, W_l's expected outcome is $U_O^l + \beta(U_h^l + \beta U_l^l) + (1-\beta)\beta U_l^l$. It is easy to see that when BW1 holds then $U_O^l + \beta(U_h^l + \beta U_l^l) + (1-\beta)\beta U_l^l > U_l^l + 2\beta U_l^l$. In that case, W_l will opt out in the first encounter if offered $\hat{s}^{W_l,3}$. However, since A_j prefers $\hat{s}^{W_h,3}$ over opting out, it will also offer $\hat{s}^{W_h,3}$ in the first encounter.

Suppose the condition BA23 of Theorem 11 holds. If the second encounter will occur then A_j will offer $\hat{s}^{W_l,3}$ and W_l will opt out with probability p_W (as defined in Lemma 10) and will accept the offer with probability $(1-p_W)$. In the third encounter, if it happens, A_j will offer $\hat{s}^{W_h,3}$ with probabilities p_A (as defined in Lemma 10) and will offer $\hat{s}^{W_l,3}$ with probability $1-p_A$. However, if the second encounter does not occur by BA2, A_j will offer $\hat{s}^{W_l,3}$ which be accepted by W_l. Therefore, W_l's expected utility if it opts out in the first encounter is: $U_O^l + \beta[(p_W(U_O^l + \beta(p_A U_h^l + (1-p_A)U_l^l))) + (1-p_W)(U_l^l + \beta U_l^l)] + (1-\beta)\beta U_l^l$. If W_l will accept $\hat{s}^{W_l,3}$ in the first encounter, its expected utility be $U_l^l + \beta U_l^l + \beta U_l^l$.

It is easy to show that since $U_l^l > U_O^l$, then $U_O^l + \beta[(p_W(U_O^l + \beta(p_A U_h^l + (1-p_A)U_l^l))) + (1-p_W)(U_l^l + \beta U_l^l)] + (1-\beta)\beta U_l^l < U_l^l + \beta U_l^l + \beta U_l^l$ holds. We can conclude that in this case W_l cannot always threaten to opt out if it is offered $\hat{s}^{W_l,3}$. □

In the rest of the section we assume that inequality BA23, holds. In such a case, as was shown in the previous lemma, the agents will use mixed strategies.

Lemma 21 (agents' behavior in the third encounter)
If the model satisfies assumptions $A0-A7$, BA2, BW1M and inequality BA23 and the agents use sequential equilibrium strategies then

1. *If W is offered $\hat{s}^{W_l,3}$ in the third encounter, W_l will accept the offer and W_h will opt out.*

2. *According to W_l's sequential equilibrium, if it is offered $\hat{s}^{W_l,3}$ in the first and second encounters then it will respectively opt out with probabilities p_0 and p_1. If the first two encounters occurred and in both of them W opted out, then A_j will choose randomly between offering $\hat{s}^{W_l,3}$ and $\hat{s}^{W_h,3}$ in the third encounter if the following equation holds:*

$$\textbf{E2A} \qquad \frac{U_l^A \phi^j p_0 p_1 + (1-\phi^j)U_O^A}{1-\phi^j+\phi^j p_0 p_1} - U_h^A = 0$$

3. *If the second encounter did not occur or A_j offered $\hat{s}^{W_h,3}$ in the second encounter and if E2A holds, then in the third encounter A_j will offer $\hat{s}^{W_l,3}$.*

Proof

1. In the last encounter there is nothing for W_l to gain from pretending to be W_h and therefore it will accept $\hat{s}^{W_l,3}$.

2. If both the first and second encounters occurred, by Lemma 19, A_j believes with probability $\frac{\phi^j p_0 p_1}{1-\phi^j+\phi^j p_0 p_1}$ that W's type is l. Therefore, by the above item, its expected utility from offering $\hat{s}^{W_l,3}$ is $U_l^A \frac{\phi^j p_0 p_1}{1-\phi^j+\phi^j p_0 p_1} + U_O^A(1-\frac{\phi^j p_0 p_1}{1-\phi^j+\phi^j p_0 p_1})$ which should be equal to U_h^A if A_j chooses randomly between them.

3. If the second encounter did not occur or A_j offered $\hat{s}^{W_h,3}$ in the second encounter, A_j's belief in the beginning of the third encounter is as in the end of the first one (by Lemma 19), i.e., $\frac{\phi^j p_0}{1-\phi^j+\phi^j p_0}$. It is easy to see that since $p_1 < 0$, $\frac{\phi^j p_0}{1-\phi^j+\phi^j p_0} > \frac{\phi^j p_0 p_1}{1-\phi^j+\phi^j p_0 p_1}$ and since $E2A$ holds, that we have $\frac{\phi^j p_0}{1-\phi^j+\phi^j p_0} U_l^A + \frac{1-\phi^j}{1-\phi^j+\phi^j p_0} U_O^A > U_h^A$ and therefore in this case A_j should offer $\hat{s}^{W_l,3}$.

$\square$

We should now consider W_l's behavior in the second encounter. If offered $\hat{s}^{W_l,3}$ it will select randomly between opting out and accepting the offer if its expected outcome is to be the same. This places restrictions on the probability q_2 in which A_j offers $\hat{s}^{W_h,3}$ in the third encounter.

Lemma 22 (A_j's probability of offering $\hat{s}^{W_h,3}$ in the third encounter)
If the model satisfies assumptions $A0 - A7$, $BA2$, $BW1M$ and inequality $BA23$ and the agents use sequential equilibrium strategies, then $q_2 = \frac{U_l^l - U_O^l}{\beta(U_h^l - U_l^l)}$.

Proof
Similar to the proof of the value of p_A in Lemma 10. $\square$

As W_l's behavior in the second encounter influences A_j's probability in offering $\hat{s}^{W_h,3}$ in the third encounter, W_l's behavior in the first encounter influences A_j's probability in the second encounter as we will explain in the next lemma.

Lemma 23 (A_j's probability of offering $\hat{s}^{W_h,3}$ in the second encounter)
If the model satisfies assumptions $A0 - A6$, $BA2$, $BW1M$ and inequality $BA23$ and the agents use sequential equilibrium strategies satisfying the properties of lemmas 19-22, then $q_1 = \frac{U_l^l - U_O^l}{\beta(U_h^l - U_l^l)}$.

Proof
If W_l opts out in the first encounter, it should consider whether the second encounter will occur or not. If so (with probability β), A_j will select randomly between offering $\hat{s}^{W_l,3}$ and $\hat{s}^{W_h,3}$. If A_j ends up offering $\hat{s}^{W_h,3}$ in the second encounter, by Lemma 21 it will offer $\hat{s}^{W_l,3}$ in the third encounter, which will be accepted by W_l. Therefore, the outcome in this case is $U_h^l + \beta U_l^l$. If A_j offers $\hat{s}^{W_l,3}$ in the second encounter, then W_l will opt out with probability p_1 and if W_l really opts out in the second encounter then by Lemma 22 in the third encounter, A_j will offer $\hat{s}^{W_h,3}$ with probability $\frac{U_l^l - U_O^l}{\beta(U_h^l - U_l^l)}$ and $\hat{s}^{W_l,3}$ otherwise. In that case, both will be accepted by W_l, and therefore its expected outcome from this case is: $(U_O^l + \beta(U_l^l \frac{U_l^l - U_O^l}{\beta(U_h^l - U_l^l)} + (1 - \frac{U_l^l - U_O^l}{\beta(U_h^l - U_l^l)})U_h^l)$. If W_l accepts $\hat{s}^{W_l,3}$ in the second encounter (with probability $1 - p_1$), it will receive a similar offer in the third encounter.

If the second encounter does not occur, but the third one does, then by Lemma 21 it will be offered $\hat{s}^{W_l,3}$ which it will accept.

Putting all the cases together we get:

$$\begin{aligned} U_O^l + \beta[q_1(U_h^l + \beta U_l^l) + (1-q_1)[(1-p_1)(U_l^l + \beta U_l^l) + \\ p_1(U_O^l + \beta(U_l^l \tfrac{U_l^l - U_O^l}{\beta(U_h^l - U_l^l)} + (1 - \tfrac{U_l^l - U_O^l}{\beta(U_h^l - U_l^l)})U_h^l))]] \\ +(1-\beta)\beta U_l^l = \\ U_l^l + 2\beta U_l^l \end{aligned}$$

Simplifying the equation we get

$$q_1 \beta U_l^l + U_l^l - U_O^l - q_1 \beta U_h^l = 0 \tag{8}$$

and we conclude that $q_1 = \frac{U_l^l - U_O^l}{\beta(U_h^l - U_l^l)}$. $\square$

The last restriction in this situation has to do with A_j's choosing randomly between $\hat{s}^{W_l,3}$ and $\hat{s}^{W_h,3}$ in the second encounter. We describe the appropriate equation in the next lemma.

Lemma 24 (A_j's behavior in the second encounter)
If the model satisfies assumptions $A0 - A7, BA2$, $BW1M$ and inequality $BA23$ and the agents use sequential equilibrium strategies, then if A_j chooses randomly between offering $\hat{s}^{W_l,3}$ and $\hat{s}^{W_h,3}$ in the second encounter the following holds:

E1A

$$\begin{aligned} &(1 - P_1 + P_1 p_1)[U_O^A + \beta(q_2 U_h^A + (1 - q2)(P_1 U_l^A + (1 - P_1)U_O^A))] \\ &+(P_1 - P_1 p_1)[U_l^A + \beta U_l^A] \\ &\quad = U_h^A(1 + \beta) \end{aligned}$$

where $P_1 = \frac{\phi^j p_0}{1-\phi^j+\phi^j p_0}$ *and* $q_2 = \frac{U_l^l - U_O^l}{\beta(U_h^l - U_l^l)}$.

Proof

In order for A_j to choose randomly between offering $\hat{s}^{W_l,3}$ and $\hat{s}^{W_h,3}$, its expected outcome from both options should be the same. A_j's outcome when it offers $\hat{s}^{W_l,3}$ is interesting. Let us denote by P_1 A_j's belief at the beginning of the second encounter after W opted out in the first encounter. According to Lemma 19, $P_1 = \frac{\phi^j p_0}{1-\phi^j+\phi^j p_0}$. If A_j will offer W $\hat{s}^{W_l,3}$ in the second encounter, W_h will opt out and W_l will opt out with probability p_1. Therefore, with probability $(1-P_1+P_1p_1)$ opting out will occur. If opting out occurs, then in the third encounter A_j will again choose randomly between offering $\hat{s}^{W_l,3}$ and $\hat{s}^{W_h,3}$. With probability q_2, A_j will offer $\hat{s}^{W_h,3}$ which will be accepted by W regardless of its type, and with probability $1-q_2$ it will offer $\hat{s}^{W_l,3}$ which will be accepted by W_l but W_h will opt out. Putting this together we get the equation

$$\begin{aligned}(1-P_1+P_1p_1)[U_O^A + \beta(q_2U_h^A + (1-q2)(P_1U_l^A + (1-P_1)U_O^A))] \\ +(P_1 - P_1p_1)[U_l^A + \beta U_l^A] \\ = U_h^A(1+\beta)\end{aligned} \tag{9}$$

Substituting the value of P_2 and the value of q_2 from Lemma 22, and simplifying the equation one can get B1A. □

In the next theorem we summarize the results in the case where there are three encounters.

Theorem 25

If the model satisfies assumptions $A0-A7, BA2,$ *BW1M and inequality BA23 and there are solutions to equations* $E1A$ *and* $E2A$ *such that* $0 < p_0 < 1$, *and* $0 < p_1 < 1$, *then there are sequential equilibrium strategies for* W_l, W_h *and* A_j *as follows:*

- *In the first encounter* A_j *offers* $\hat{s}^{W_l,3}$, W_h *will opt out and* W_l *will choose to opt out with probability* p_0 *and will choose to accept the offer with probability* $1-p_0$.
- *In the second encounter (if it occurs)* A_j *will offer* $\hat{s}^{W_l,3}$ *with probability* $q_1 = \frac{U_l^l - U_O^l}{\beta(U_h^l - U_l^l)}$ *and will offer* $\hat{s}^{W_h,3}$ *with probability* $1-q_1$. *If* $\hat{s}^{W_h,3}$ *is offered,* W *will accept the offer (regardless of its type), and if* $\hat{s}^{W_l,3}$ *is offered,* W_h *will opt out and* W_l *will will accept the offer with probability* $1-p_1$ *and will opt out with probability* p_1.

- *In the third encounter (if it happens), W_l will always accept $\hat{s}^{W_l,3}$ or $\hat{s}^{W_h,3}$, and W_h will accept $\hat{s}^{W_h,3}$ but will opt out if it is offered $\hat{s}^{W_l,3}$.*

  If the second encounter did not occur or A_j offered $\hat{s}^{W_h,3}$, then in the third encounter A_j will offer $\hat{s}^{W_l,3}$. Otherwise, it will offer $\hat{s}^{W_h,3}$ with probability $q_2 = \frac{U_l^l - U_O^l}{\beta(U_h^l - U_l^l)}$ and with probability $1 - q_2$ it will offer $\hat{s}^{W_l,3}$.

Proof
Clear from the above lemmas. □

In order to have appropriate solutions to equations $E1A$ and $E2A$ that can serve as probabilities (i.e., solutions that are between 0 and 1), then $\frac{(1-\phi^j)(U_l^A - U_O^A)}{\phi^j(U_l^A - U_h^A)} < 1$ must hold. We will demonstrate this case in our communication systems example.

Example 8
Suppose the situation is as in Example 2 but with the following specification: $U_l^A = 20, U_h^A = 10, U_O^A = 8, U_h^l = 5, U_l^l = 2$, and $U_O^l = 1.5$. Suppose that there are three possible encounters where the probability of the second and third one is $\beta = 0.75$ and $\phi^j = .75$. In such situations $q_1 = q_2 = .2222222222$, $p_0 = .443809$ and $p_1 = 0.1502145$.

6 Extension of the Model

There are several possible extensions of the model. Here we discuss the cases where there are many resources in the environment, and situations where there are more than two types of agents.

6.1 Many resources

Suppose that there are several resources in the environment and at any given time, only two agents may share the same resource, after the agents have reached a detailed agreement. There may be two types of resources: available ones, and resources that are already in use by other agents.[34] In such an environment, when an agent needs a resource, it may check if there is such a resource that is not in use. However, if all the resources of the type that is needed are already in use, it may find the resources that are being used by only one agent, and based on its beliefs about their types and its utility for using the specific resources, it can

[34] The agents may use some *Test and Set* mechanism. This will prevent the situation in which two agents would like to get access at exactly the same time.

decide with whom to start the negotiations. We assume that an agent cannot negotiate with more than one agent at a time.

Example 9
We return to the example of the communication systems. Suppose there are two public communication lines available for payment; one is used by BankA and the other by BankC. BankB is experiencing an exceptionally high workload as in the previous examples. BankB then needs to reach an agreement with one of the other banks on sharing the public communication line it is using.

Let us assume that there is only one encounter, and we denote by W the agent that is waiting for a resource. For any resource R and an agent A^R that uses it, W has some probabilistic belief about its type and about the belief of A^R about W's type. For each of these types, W computes the possible outcome of the negotiation with A^R_j where $j \in \mathbf{Type}$ and computes the expected utility for W from it (denoted by $U(A^R_j), j \in \mathbf{Type}$). Using its own beliefs about the type of its opponent, W computes the overall expected utility from R.

After computing the expected utility for all the resources, W chooses the one with the highest expected utility and negotiates according to the strategies of the previous section. After choosing a resource, it is easy to prove that the agent may not change its decision, i.e., it will not stop negotiating with one agent, and start a new negotiation process with another agent about a different resource.

6.2 More than two types of agents

Suppose there are more than two types of agents (i.e., $k > 2$). The situation is similar to when there are only two, but the agents need to take more options into consideration. If there is only one encounter and A_j offers $\hat{s}^{W_r,3}$ in the second iteration, then if W' type is $i \leq r$ it will accept the offer. If $i > r$, W_i will opt out. Suppose that the maximum expected utility for A_j in such a case is from $\hat{s}^{W_r,3}$ and that it isn't worthwhile for A_j to offer less in order to gain information. If so then A_j will offer $\hat{s}^{W_r,3}$ and if it is accepted, A_j will know that W's type is at most r and will update its belief accordingly using Bayes's rules. It is easy to prove that in such a case A_j will offer $\hat{s}^{W_r,3}$ in the second encounter as well. However, if W opts out, A_j may conclude that W's type is greater than r and update its belief. In such a case, A_j will offer in the next encounter $\hat{s}^{W_x,3}$ for some $x > r$. The question, as in the case of two types, is if W's type is less than r whether it is worthwhile for W to opt out or whether there is an equilibrium of pure strategies, or whether the agents should use mixed strategies. In the case where there are no pure strategies, the process of identifying the probabilities of the mixed strategies is similar to the case where there are only two types of agents. However, it will require solving more equations.

7 Complexity and Implementation Issues

We have constructed a library of meta-strategies that are appropriate for negotiation situations characterized by time with one encounter. When an automated negotiator participates in a specific negotiation environment, it could search for the appropriate meta-strategy in the library, initiate some variables and use it in the negotiation. It is assumed that agents believe that they will all assess the situation the same way, and will use the equilibrium strategies, either by finding it in a similar meta-strategies library or by computing it in other ways. We will report on the implementation issues in a different paper [39], but we discuss here some important related questions.

The library is kept in an OR/AND tree where the internal nodes consist of conditions and the meta-strategies are stored in the leaves of the tree. The number of negotiators is an example of a condition. When searching for the appropriate meta-strategy, if any exists, (for example, where there are two negotiators in the environment), the agent will continue its search in the subtree that consists of the meta-strategies for bilateral negotiation. Other examples of simple conditions are whether the agents can opt out from the negotiations or not, whether $c_A > c_W$, etc.

A meta-strategy that is stored in the leaves consists of a compiled function and a set of names of variables. Some of the variables are instantiated during the search in the tree (e.g., c_A) and some are instantiated during the negotiation when the agents use the function to decide what to do next (e.g., t the iteration of the negotiation).

There are procedures that help a designer of an agent to add new meta-strategies to the library, and we intend to add multiple-encounter strategies found in this paper to that library.

Concerning the motivation for constructing the library, there are two approaches to finding equilibria in incomplete information models. One is the straight game theory approach: search for Nash or sequential strategies. The other is the economist's standard approach: set up a maximization problem and solve it by using calculus [51]. 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., [59]). However, when applied to situations such as ours, the maximization technique is less appropriate since 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 and there is no general way to make the initial guess.

In situations of multistage negotiation (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. It is often best to work through problems like this backward.

In our negotiation protocol there are $\frac{M^{|\mathbf{Agent}|-1}}{(|\mathbf{Agent}|-1)!}$ possible actions in the first period of each iteration (i.e., the number of possible agreements) and $3^{|\mathbf{Agent}|}$ possible combinations of actions in the second time period of each iteration. 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 than agreement), then the overall number of pure strategies is $O((\frac{M^{|\mathbf{Agent}|-1}}{(|\mathbf{Agent}|-1)!})^{\hat{T}})$.

However, since we consider cases of incomplete information, 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. This makes it too time consuming to compute the strategies in real time. Therefore, we suggest that finding the equilibrium strategies be done *before* the negotiation process starts. In our papers, we present appropriate strategies to varied situations. The situations are characterized by several factors of the environments (e.g., number of agents, purpose of the negotiation) and the agents' utility functions. That is why, in this paper, we took the step of calculating the exact conditions in which each of the sequential-equilibrium strategies are applicable. These conditions will appear in the internal nodes on the OR/AND meta-strategies tree.

Another important question is how the autonomous agents will initiate their beliefs. One possibility is that they will have some general information about the distribution of the agent types among their opponents. For example, it may be known to BankA that half of banks are of type l and half are of type h. Of course, in such situations, if BankA does not have any additional information, it will believe with probability $\frac{1}{2}$ that its opponent is of type l. During the negotiation encounters it may update its belief.

Techniques that were presented by Bacchus et al. [3], for assigning degrees of beliefs by an intelligent agent based on known facts, including statistical knowledge, can be used in our situation.

If there is no prior information about the opponent's type, then the agent can always assume that there is equal distribution of the types of agents and use it for its prior probability beliefs.

8 Conclusion and Open Questions

In this paper we presented a strategic model of negotiation that takes the passage of time into consideration, addresses situations where agents may negotiate more than once with each other, and where they have incomplete information about one another. Our results satisfy the desired properties listed in Section 1.2.

Distribution: 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. The agents' types do not play any role in the negotiation protocol. They only influence the negotiation strategies.

Instantaneously : The negotiation will always end in the second iteration.

Conflicts are Avoided : In some of the cases, the incomplete information introduces inefficiency: one of the agents may opt out of the negotiation. In the case of two encounters, agents will always reach an agreement when BA1 and BA1.1 hold or when BA1, BA1.2 and BW1 hold.
If W's type is h then BA1, BA1.2 or BA2 hold, and if BW2 holds, opting out will occur in the first encounter, however, as we explained above, BW2 rarely holds.
If BA2 and BW1 hold and if the reverse of inequality BA23 holds, then an agreement will always be reached in the first encounter, but if W's type is h, opting out will occur in the second encounter. If inequality BA23 holds, then there is a high probability that opting out will occur in the first encounter, and lower probability that it will occur in the second one.
Similar events occur when there are more than two encounters.

Efficiency : The resource is not in use only when there is no agent in the group that currently needs the resource. However, opting out may introduce inefficiency to the system.

Simplicity : Given specific specification of an environment, the strategies are simple: the agent's action depends only on the current situation and its beliefs. The offers and the communications are simple.
However, as was demonstrated in the paper, computing the strategies is sometimes a very difficult task and it is not recommended that this be done on-line. Therefore, we have investigated a large range of situations, and describe the appropriate strategies for them. If the agents participate in such situations, their designers provide them with these strategies.

Stability : In the situations that we have considered we have found sequential equilibrium strategies, either with pure or mixed strategies.

Symmetry : The coordination mechanism we presented does not treat agents differently because of non-relevant attributes.

We believe that our model can be useful in other situations beside the ones we analyzed in the paper. In this paper we consider the problem of resource allocation, but the dual problem of task distribution can also be analyzed using this model [34]. There are several assumptions which can be relaxed to make the model appropriate to more realistic domains. We list some of the possible extensions here.

- Negotiation on multiple attributes.
- The agents may collect information on one agent, while negotiating with another one.
- Negotiations where A's type plays an important role.
- The occurrence of one encounter may influence the probability of future ones.

We leave this for on-going work.

References

[1] R. J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1):67–96, 1974.

[2] L. M. Ausubel and R. J. Deneckere. Stationary sequential equilibria in bargaining with two sided incomplete information. Discussion Paper 784, Center for Mathematical Studies in Economics and Management Science, Northwestern University, 1988.

[3] F. Bacchus, A. Grove, J. Halpern, and D. Koller. From statistics to beliefs. In *Proc. of AAAI-92*, pages 602–608, California, 1992.

[4] 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.

[5] E. Bond and L. Samuelson. Durable good monopolies with rational expectation and replacement sales. *Rand Journal of Economics*, 15:336–345, 1984.

[6] E. Bond and L. Samuelson. The Coase conjecture need not hold for durable good monopolies with depreciation. *Economics Letters*, 24:93–97, 1987.

[7] 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.

[8] K. Chatterjee and L. Samuelson. Bargaining with two-sided incomplete information: An infinite horizon model with alternating offers. *Review of Economic Studies*, 54:175–192, 1987.

[9] I. K. Cho. Characterization of stationary equilibria in bargaining models with incomplete information. Unpublished paper, Department of Economics, University of Chicago, 1989.

[10] P. Cohen and H. Levesque. Teamwork. *Noûs*, pages 487–512, 1991.

[11] 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.

[12] 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.

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

[14] J. Doyle. Some theories of reasoned assumptions: an essay in rational psychology. Technical Report 83-125, Department of Computer Science, Carnegie Mellon University, Pittsburgh, PA, 1983.

[15] J. Doyle. Rational belief revision. Technical report, MIT, Artificial Intelligence Laboratory, Cambridge, Massachusettes, 1990. Unpublished paper.

[16] D. Dubois and H. Prade. A survey of belief revision and updating rules in various uncertainty models. *International Journal of Intelligent Systems*, 9:61–100, 1994.

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

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

[19] E. Ephrati and J. S. Rosenschein. Reaching agreement through partial revelation of preferences. In *Proceedings of the Tenth European Conference on Artificial Intelligence*, pages 229–233, Vienna, Austria, August 1992.

[20] R. Fagin, G. Kuper, J. Ullman, and M. Vardi. Updating logical databases. In *Advances in Computing Research*, volume 3, pages 1–18, 1986.

[21] D. Fudenberg and J. Tirole. *Game Theory*, chapter 8. MIT Press, Cambridge, Ma, 1991.

[22] P. Gardenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States.* MIT Press, 1988.

[23] J. Y. Halpern and Yoram Moses. Knowledge and common knowledge in a distributed environment. *Journal of the Association for Computing Machinery*, 37(3), 1990.

[24] G. Harman. *Change in View.* MIT Press, Cambridge, MA, 1986.

[25] W. L. Harper. Rational belief change, popper functions, and counterfactuals. In W. L. Harper and C. A. Hooker, editors, *Foundations of Probability Theory, Statistical Inference, and Statistical Theories of Science*, volume 1, pages 73–115. Reidel, Dordrecht, 1976.

[26] O. D. Hart and J. Tirole. Contract renegotiation and Coasian dynamics. *Review of Economic Studies*, pages 509–540, 1988.

[27] H. Katsuno and A. Mendelzon. Knowledge Base Revision and Minimal Change. *Artificial Intelligence*, 52:263–294, 1991.

[28] Hirofumi Katsuno and Alberto Mendelzon. On the difference between updating a knowledge base and revising it. In *Proc. of KR91*, pages 387–394, 1991.

[29] S. Kraus and D. Lehmann. Knowledge, belief and time. *Theoretical Computer Science*, 58:155–174, 1988.

[30] S. Kraus and D. Lehmann. Designing and building a negotiating automated agent. *Computational Intelligence*, 11(1):132–171, 1995.

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

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

[33] S. Kraus and J. Wilkenfeld. A strategic negotiations model with applications to an international crisis. *IEEE Transaction on Systems Man and Cybernetics*, 23(1):313—323, 1993.

[34] S. Kraus, J. Wilkenfeld, and G. Zlotkin. Multiagent negotiation under time constraints. *Artificial Intelligence*, 75(2):297–345, 1995.

[35] D. Kreps and R. Wilson. Reputation and imperfect information. *Journal of Economic Theory*, 27:253–279, 1982.

[36] D. Kreps and R. Wilson. Sequential equilibria. *Econometrica*, 50:863–894, 1982.
[37] K. Kuwabara and V. Lesser. Extended protocol for multistage negotiation. In *Proc. of the Ninth Workshop on Distributed Artificiall Intelligence*, pages 129–161, 1989.
[38] B. Laasri, H. Laasri, and V. Lesser. A generic model for intelligent negotiating agents. *International Journal on Intelligent Cooperative Information Systems*, 1(2):291–317, 1992.
[39] G. Lemel. A strategic model for negotiation among autonomous agents. M.Sc. thesis, Dept. of Mathematics and Computer Science, Bar-Ilan University, Ramat Gan (written largely in Hebrew)., 1995.
[40] V. R. Lesser, J. Pavlin, and E. H. Durfee. Approximate processing in real time problem solving. *AI Magazine*, 9(1):49–61, 1988.
[41] V.R. Lesser. A retrospective view of FA/C distributed problem solving. *IEEE Transactions on Systems, Man, and Cybernetics*, 21(6):1347–1362, 1991.
[42] D. Lewis. *Counterfactuals*. Blackwell, Oxford, 1973.
[43] R. D. Luce and H. Raiffa. *Games and Decisions*. John Wiley and Sons, 1957.
[44] V. Madrigal, T. Tan, and R. Werlang. Support restrictions and sequential equililibria. *Journal of Economic Theory*, 43:329–334, 1987.
[45] 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.
[46] P. Milgrom and J. Roberts. Predation, reputation, and entry deterrence. *Journal of Economic Theory*, 27:280–312, 1982.
[47] T. Moehlman, V. Lesser, and B. Buteau. Decentralized negotiation: An approach to the distributed planning problem. *Group Decision and Negotiation*, 2:161–191, 1992.
[48] M. J. Osborne and A. Rubinstein. *Bargaining and Markets*. Academic Press Inc., San Diego, California, 1990.
[49] W. V. Quine and J. S. Ullian. *The Web of Belief*. Random House, New York, 1978. Second Edition.
[50] A. S. Rao and N. Y. Foo. Formal theories of belief revision. In *Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning*, pages 369–380. Morgan Kaufmann, May 1989.
[51] E. Rasmusen. *Games and Information*. Basil Blackwell Ltd., Cambridge, MA, 1989.
[52] J. S. Rosenschein. *Rational Interaction: Cooperation Among Intelligent Agents*. PhD thesis, Stanford University, 1986.
[53] J. S. Rosenschein and G. Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiation Among Computers*. MIT Press, Boston, 1994.
[54] A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–109, 1982.
[55] A. Rubinstein. A bargaining model with incomplete information about preferences. *Econometrica*, 53(5):1151–1172, 1985.
[56] T. Sandholm. An implementation of the contract net protocol based on marginal cost calculations. In *Proc. of AAAI-93*, pages 256–262, 1993.
[57] L. J. Savage. *The Foundations of Statistics*. Dover Publications, New York, 1972. Second edition.
[58] R.G. Smith and R. Davis. Negotiation as a metaphor for distributed problem solving. *Artificial Intelligence*, 20:63–109, 1983.
[59] W. Spivey and R. Thrall. *Linear Optimization*. Holt, Rinehart and Winston, 1970.
[60] R. C. Stalnaker. *Inquiry*. MIT Press, Cambridge, MA, 1984.
[61] K. P. Sycara. Persuasive argumentation in negotiation. *Theory and Decision*, 28:203–242, 1990.
[62] K.P. Sycara. *Resolving Adversarial Conflicts: An Approach to Integrating Case-Based and Analytic Methods*. PhD thesis, School of Information and Computer Science, Georgia Insti-

tute of Technology, 1987.

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

[64] M. Winslett. Is belief revision harder than you thought? In *Proceedings of AAAI-86*, pages 421–427, Philadelphia, 1986.

[65] G. Zlotkin and J. S. 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.

[66] G. Zlotkin and J. S. Rosenschein. Incomplete information and deception in multi-agent negotiation. In *Proc. IJCAI-91*, pages 225–231, Australia, 1991.

[67] 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.