Alternating-Offers Bargaining under One-Sided Uncertainty on Deadlines

Francesco Di Giunta and **Nicola Gatti**[1]

Abstract. Alternating-offers is the most prominent negotiation protocol for automatic bilateral bargaining. Nevertheless, in most settings it is still not known how two fully rational agents should behave in the protocol. In this paper we study the finite-horizon alternating-offers protocol under one-sided uncertain deadlines. We make a novel use of backward induction in studying bargaining with uncertainty; we employ a "natural" system of beliefs and find, when it exists, the pertinent pure strategy sequential equilibrium. We further show, as an intrinsic limitation of the protocol, that for some parameter values there is no pure strategy sequential equilibrium, whatever system of beliefs is employed.

1 Introduction

Automated negotiation is a prominent research area of distributed artificial intelligence which studies the processes whereby rational software agents try to solve disputes and reach mutually beneficial agreements [7]. It is well known that the automation of the negotiating agents leads to more effective negotiations since agents are more efficient than humans in finding optimal agreements [14].

Among the negotiation settings for commercial transactions, a very common one is *bargaining*: a buyer and a seller try to agree on the choice of the value of some variables, named *issues*, that define the transaction they are carrying out (e.g. the price and quantity of a traded good). The formalized study of bargaining is commonly carried out with game-theoretical tools [9] in which one distinguishes the negotiation *protocol* and the negotiation *strategies*: the protocol sets the negotiation rules, specifying which actions are allowed and when [10]; the strategy defines an agent's possible specific behavior in the negotiation. Given a protocol, rational agents should employ strategies that maximize their payoffs, and the classic game-theoretic approach prescribes that agents employ equilibrium strategies, where the notion of equilibrium is Nash equilibrium or extensions [5].

The best known protocol for bilateral bargaining, the *alternating-offers* protocol, pioneered by Ståhl [16], has reached an outstanding place in literature thanks to Rubinstein [11], and comes in many variations. Basically, an agent starts by offering values for the issues under dispute to her opponent, who can accept or make a counteroffer or exit the negotiation. If a counteroffer is made, the process is repeated until one of the agents accepts or exits the negotiation. There can be one or more bargained issues, leading respectively to *one-issue* and *multi-issue* bargaining. The issues are usually continuous, i.e. they are real-valued variables. The agents have reservation values, temporal utility discount factors, and negotiation deadlines. If the agents have finite negotiation deadlines, the protocol is a *finite-horizon* one; otherwise, it is *infinite-horizon*.

[1] Politecnico di Milano,Italy, email: {digiunta, ngatti}@elet.polimi.it

Although much economics and computer science literature deals with the alternating-offers protocol (see e.g. [7]), several major problems are still open, the main ones concerning *multi-issue* and *incomplete information* bargaining. Easy and general solutions are currently available when the bargaining is on one issue and every pertinent information is common knowledge between the two agents, but both assumptions are very restrictive. The problem of efficient multi-issue bargaining has been addressed in [2]. The problem of incomplete knowledge is harder. Rubinstein [12] has provided the first results about uncertainty over two possible time discount factors of one of the two agents. Several other authors have followed his research line, providing results on very narrow lacks of information completeness. We refer to [1] for a survey.

Given the difficulties in finding classic solutions to alternating-offers problems with incomplete information, several authors have proposed non-standard approaches. The best known proposal is Fatima, Wooldridge, and Jennings's framework in [4]. Their approach is based on the *negotiation decision functions* paradigm [3], where the agents are supposed to constraint themselves to employ one of some predefined bidding tactics in their bargaining.

The aim of our paper is to analyze and solve the finite-horizon alternating-offers protocol under incomplete information on the bargaining deadline of one of the two agents (*one-sided* incomplete information). We adopt a classical approach, where the agents are not *a priori* assumed self-restricted to any class of strategies, in the belief that this approach is more coherent with the multiagent system paradigm (see, e.g., the discussion in [15]). Furthermore, in accordance with virtually all the literature on the topic, we employ pure strategies rather than mixed ones.

Our original contributions are: (*1*) the solution of the single-issue finite-horizon alternating-offers bargaining under one-sided uncertainty on deadlines, for a wide range of bargaining parameters; (*2*) the proof of non-existence of a general solution in pure strategies.

The paper is structured as follows: in the next section we review the finite-horizon alternating-offers protocol with complete information and its well known solution, in order to introduce basic concepts and state notation; in Section 3 we analyze the protocol under one-sided uncertain deadlines; in Section 4 we show that the general problem has no solution in pure strategies; Section 5 summarizes our conclusions.

2 Complete information alternating-offers

The protocol we study is a discrete time finite-horizon alternating-offers bargaining protocol on one continuous issue (say, a price).[2]

[2] The multi-issue problem [2] is orthogonal to the incomplete information problem. So, for the sake of simplicity, we consider one-issue bargaining.

Formally, the buyer b and the seller s can act at times $t \in \mathbb{N}$. The *player function* $\iota : \mathbb{N} \rightarrow \{b, s\}$ returns the agent that acts at time t and is such that $\iota(t) \neq \iota(t+1)$. Possible actions $\sigma^t_{\iota(t)}$ of agent $\iota(t)$ at time $t > 0$ are:

(*1*) *offer*$(\overline{x})$, where $\overline{x} \in \mathbb{R}$,
(*2*) *exit*,
(*3*) *accept*.

At $t = 0$ the only allowed actions are (*1*) and (*2*). If $\sigma^t_{\iota(t)} = accept$ the bargaining stops and the *outcome* is $(\overline{x}, t)$, where $\overline{x}$ is the tuple such that $\sigma^{t-1}_{\iota(t-1)} = offer(\overline{x})$. If $\sigma^t_{\iota(t)} = exit$ the bargaining stops and the outcome is $NoAgreement$. Otherwise the bargaining continues to the next time point.

Each agent i has an utility function $U_i : (\mathbb{R} \times \mathbb{N}) \cup \{NoAgreement\} \rightarrow \mathbb{R}$, that represents her gain on the possible bargaining outcomes. Each utility function U_i depends on three parameters of agent i:

- the *reservation price* $RP_i \in \mathbb{R}^+$,
- the *temporal discount factor* $\delta_i \in (0, 1]$,
- the *deadline* $T_i \in \mathbb{N}$, $T_i > 0$.

Exactly, if the outcome of the bargaining is an agreement (x, t), then the utility functions U_b and U_s are respectively:

$$U_b(x,t) = \begin{cases} (RP_b - x)\delta_b^t & \text{if } t \leq T_b \\ -1 & \text{otherwise} \end{cases}, U_s(x,t) = \begin{cases} (x - RP_s)\delta_s^t & \text{if } t \leq T_s \\ -1 & \text{otherwise} \end{cases}.$$

If the outcome is $NoAgreement$, then $U_b(NoAgreement) = U_s(NoAgreement) = 0$.[3]

Some standard hypothesis are assumed. *Feasibility*: $RP_b \geq RP_s$. *Rationality*: it is common knowledge that each agent will act to maximize her utility. *Benevolence*: it is common knowledge that if an agent has to choose between two outcomes which are indifferent for her but not for her opponent, she will choose the one that is better for her opponent.

The bargaining strategies that the agents should select also depend on the knowledge that the agents have of the protocol and of each other's utility function. In *complete information* bargaining it is assumed that the protocol and the utility functions (including the values of RP_i, δ_i and T_i) are common knowledge between the two agents.

The appropriate notion of solution for a complete information extensive form game like the one we are dealing with is *subgame perfect equilibrium* [6]. Subgame perfect equilibrium strategies can be easily found by *backward induction*, as we discuss in the following.

No agent is willing to bargain after her deadline, when any agreement would have negative utility. Therefore, at time $\overline{T} = \min\{T_b, T_s\}$ the acting agent – let's say s – would accept any offer with non-negative utility. Therefore, at time $\overline{T} - 1$ her opponent b could safely offer RP_s (which would be accepted) or could accept any possible previous offer x which is not worse than offering RP_s (i.e., $U_b(x, \overline{T} - 1) \geq U_b(RP_s, \overline{T})$). Therefore, at time $\overline{T} - 2$ agent s could safely offer the maximum x such that $U_b(x, \overline{T} - 1) \geq U_b(RP_s, \overline{T})$, i.e. x such that $U_b(x, \overline{T} - 1) = U_b(RP_s, \overline{T})$, or accept any possible previous offer which is not worse than offering x. This reasoning can be inductively carried on until the beginning of the game, finding an offer that agent $\iota(0)$ would do and her opponent would accept.

At each time point t, from $\overline{T}$ back, it is therefore possible to know which offer would be made by agent $\iota(t)$ if she would make an offer: we denote this series of offers by $x^*(t)$. In order to provide a recursive formula for $x^*(t)$, we introduce the notion of backward propagation: given value x and agent i, we call *backward propagation* of value x for agent i the value y such that $U_i(y, t-1) = U_i(x, t)$; we will employ the arrow notation $x_{\leftarrow i}$ for backward propagations.[4] The calculation of $x^*(t)$ can now be stated like this:

$$x^*(t-1) = \begin{cases} RP_{\iota(t)} & \text{if } t = \overline{T} \\ (x^*(t))_{\leftarrow \iota(t)} & \text{if } t < \overline{T} \end{cases},$$

and is computationally linear with $\overline{T}$. A backward induction and the relevant $x^*(t)$ values can be plotted on the space (x, t) as in Figure 1.

The backward induction reasoning so far sketched proves the following result [9]:

Proposition 2.1 *Finite-horizon alternating-offers over single-issue bargaining with complete information has one and only one subgame perfect equilibrium. The equilibrium strategies for* $t < \overline{T}$ *are*

$$\sigma^t_{\iota(t)} = \begin{cases} accept \quad \text{if } \begin{cases} t > 0 \\ \sigma^{t-1}_{\iota(t-1)} = offer(x) \text{ with } U_{\iota(t)}(x, t) \geq U_{\iota(t)}(x^*(t), t+1) \end{cases} \\ offer(x^*(t)) \quad \text{otherwise} \end{cases}$$

and when $t = \overline{T}$

$$\sigma^t_{\iota(\overline{T})} = \begin{cases} accept \quad \text{if } \sigma^{\overline{T}-1}_{\iota(\overline{T}-1)} = offer(x) \text{ with } U_{\iota(\overline{T})}(x, \overline{T}) \geq 0 \\ exit \quad \text{otherwise} \end{cases}.$$

The agreement is therefore achieved at time $t = 1$ *on the price* $x^*(0)$.

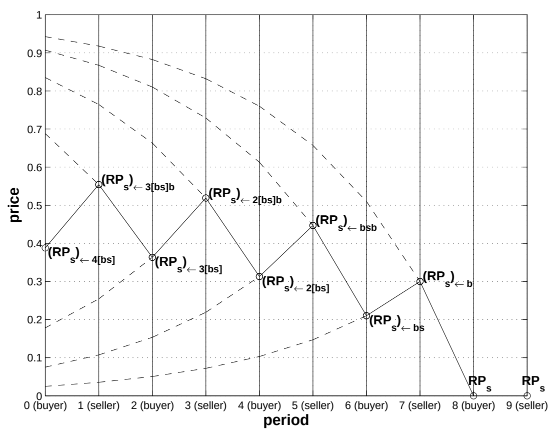

Figure 1. Backward induction with $RP_b = 1$, $RP_s = 0$, $\delta_b = 0.7$, $\delta_s = 0.7$, $T_b = 9$, $T_s = 10$, $\iota(0) = b$

3 Alternating-Offers with Uncertain Deadlines

Let us consider one-issue alternating-offers bargaining when one of the two agents does not exactly know her opponent's deadline. We will assume that the uncertain deadline is the buyer's.

As is customary in game theory in order to avoid underdetermined problems, we assume that b's possible deadlines are distributed according to a probability distribution on $\mathbb{R}^+$ which is common knowledge between the agents. We further assume that the support of b's deadline distribution is bounded; and since the agents

[3] Notice that employment of $U_i = -1$ after agent i's deadline allows a rational behaviour of the agent after her deadline: an agent prefers to exit the negotiation than to reach any agreement.

[4] If a value x is backward propagated n times along the level curves of agent i, we write $x_{\leftarrow n[i]}$. If a value is backward propagated along the level curves of more than one agent, we list them left to right in the subscript; for instance, $x_{\leftarrow i3[j]}$ is value x backward propagated along the level curves of agent i and subsequently three times along the curves of agent j.

can act only at times $t \in \mathbb{N}$, we can assume, without lost of generality, that b's deadline has a finite distribution on $\mathbb{N}$. We denote by $\mathcal{T}_b = \{T_{b_1}, \ldots, T_{b_m}\}$ the set of possible deadlines of b, by $\mathcal{P}_b^0 = \{\alpha_{b_1}^0, \ldots, \alpha_{b_m}^0\}$ the pertinent probability distribution, and by BT_b^0 the couple $BT_b^0 = \langle \mathcal{T}_b, \mathcal{P}_b^0 \rangle$. Agent b's real deadline is known only to b itself: it is her *private information*.

In games with uncertainty, rational agents try to maximize their expected utilities relying on their beliefs about the opponent's private information and such beliefs are updated during the game, depending on which actions have been actually undertaken. The set of beliefs held by each agent over the other's private information after every possible sequence of actions in the game is called a *system of beliefs* and is usually denoted by μ. This beliefs are probabilistic and their values at time $t = 0$ is a given problem data. How beliefs evolve during the game, instead, is part of the solution which should be found for the game. A solution to an incomplete information bargaining is therefore a suitable couple $a = \langle \mu, \sigma \rangle$ called *assessment*.

An assessment $a = \langle \mu, \sigma \rangle$ must be such that the strategies in σ are mutual best responses given the probabilistic beliefs in μ (*rationality*); and the beliefs in μ must reasonably depend on the actions prescribed by σ (*consistency*). Different notions of solution differ on how they specify these two requirements.

We will employ the most common notion of solution for incomplete information bargaining, namely the the *sequential equilibrium* of Kreps and Wilson [8]. For a sequential equilibrium $a^* = \langle \mu^*, \sigma^* \rangle$, with $\sigma^* = \langle \sigma_b^*, \sigma_s^* \rangle$, the rationality requirement is specified as *sequential rationality*. Informally, after every possible sequence of actions S, on or off the equilibrium path, the strategy σ_s^* must maximize s's expected utility given s's beliefs prescribed by μ for S, and given that b will there on act according to σ_b^*; and *vice versa*. The notion of consistency is defined as follows: assessment a is *consistent in the sense of Kreps and Wilson* (or simply *consistent*) if there exists a sequence a_n of assessments, each with completely mixed strategies and such that the beliefs are updated according to Bayes rule, that converges to a.

The method we will employ to find sequential equilibria is: (*1*) *a priori* fix a reasonable system of beliefs $\overline{\mu}$; (*2*) use backward induction to find, if they exist, the strategies σ of sequentially rational assessments a with beliefs $\overline{\mu}$; (*3*) *a posteriori* prove the consistency of the assessment.

We will use a system of beliefs $\overline{\mu}$ which is easy and natural: after any sequence of actions, agent s just excludes those deadlines T_{b_h}, among the initially possible ones, that have already expired and normalizes the probabilities of the future ones. Thus, in particular, the beliefs about the deadlines after a sequence of actions S depend only on the length of S. For it will be useful later, we introduce the *deadline function* $d(t)$, whose value is the probability, given at time t according to $\overline{\mu}$, that time t itself is a deadline for agent b.

Given the system of beliefs $\overline{\mu}$, we can find the sequentially rational strategies by backward induction. But the use of backward induction in this context is more involved than in the complete information bargaining and requires some explanations.

With complete information the backward induction can start at the earlier of the two deadlines; but with our incomplete information framework, the earlier deadline is not *a priori* known. Nevertheless, the backward induction can start at $\overline{T} = \min\{\max_h\{T_{b_h}\}, T_s\}\}$, because it is *a priori* known that after time $\overline{T}$ agent b will exit the negotiation; if the bargaining process would reach time $\overline{T}$, then agent gent $\iota(\overline{T})$ would accept any nonnegative utility offer; therefore, at time $\overline{T} - 1$ agent $\iota(\overline{T} - 1)$, if she would make an offer, would offer $\iota(\overline{T})$'s reservation price.

It is therefore possible, like in the complete information setting of Section 2, to backward infer the offers $x^*(t)$ that the agents would do if they would choose to make an offer. But there are some complications in the construction of the optimal offers $x^*(t)$.

In the complete information case the values $x^*(t)$ are (i) the optimal offers, (ii) the values to be backward propagated, and (iii) the backward propagated values. In incomplete information bargaining, instead, this three series of values are distinct in general. Let us see these topics in more detail.

First, in complete information bargaining the optimal offer $x^*(t)$ is simply the backward propagation of $x^*(t+1)$; with incomplete information this is not generally true: the backward propagated value is only the best among the surely accepted offers. But if at a time t, such that $\iota(t) = s$, there is a high probability that time $t+1$ is a deadline of b, then agent s could prefer to offer RP_b rather than offer the best among the surely accepted offers, although offer RP_b could be rejected (if $t+1$ is not b's real deadline).

Second, as optimal offers of agent $s = \iota(t+1)$ could be rejected, the value that should be backward propagated from time $t+1$ to time t, in order to obtain the best one among the offers of b at time t that would be surely accepted by s, is not $x^*(t+1)$ in general. The right value to be backward propagated is the value that we call *equivalent value* and is defined as follows: given an offer x of s, the equivalent value of x, denoted by $e(t)$, is the value such that $U_s(e(t), 0) = EU_s(x)$.[5] If offered value $(e^*(t+1))_{\leftarrow s}$ at time t, agent s would accept it instead of offering $x^*(t+1)$; if offered a worse value, she would refuse it and counteroffer $x^*(t+1)$.

Summarily, backward propagation should be applied in general to equivalent values e^* rather than to best offers x^*, and best offers x^* of agent s are not always the backward propagated values. Formulas to find equivalent values and best offers are easy to find. If $\iota(t) = b$ then $e^*(t) = x^*(t)$ and $x^*(t) = (e^*(t+1))_{\leftarrow s}$ Conversely, if $\iota(t) = s$ then $e^*(t) = RP_s + \delta^{t-1} \cdot EU_s(x^*(t))$ and $x^*(t) = \arg \max_{x = RP_b, (e^*(t+1))_{\leftarrow b}} \{EU_s(x)\}$ where:

$$
\begin{aligned}
EU_s(RP_b) &= (1 - d(t)) \cdot \Big(d(t+1) \cdot U_s(RP_b, t+1) + (1 - d(t+1)) \cdot \\
&\qquad \cdot U_s(x^*(t+1), t+2) \Big) \\
EU_s((e^*(t+1))_{\leftarrow b}) &= (1 - d(t)) \cdot U_b((e^*(t+1))_{\leftarrow b}, t+1)
\end{aligned}
$$

Given $x^*(\cdot)$, the optimal strategy σ_i^* for agent i at time t is: if t is i's deadline and at $t-1$ she received an offer that gives her negative utility, then exit; if t is i's deadline and at $t-1$ she received an offer that gives her nonnegative utility, then accept; if t is not i's deadline and at time $t-1$ she received an offer not worse for her than $(x^*(t))_{\leftarrow i}$, then accept; otherwise offer $x^*(t)$.

The above backward construction provides the sequentially rational strategies with the system of beliefs $\overline{\mu}$ we have chosen, when they exist. But actually, they could not exist with some parameter setting. This can happen because one agent (precisely, the buyer) might deviate from strategy σ^* to offer something unacceptable for the opponent in order to be refused and to later be in a much stronger position and gain more. An example is given in Fig. 2. Here the backward induction construction imposes that at time 6 an offering buyer should make offer $e_{\leftarrow s}^7$; but it is easily seen that it is more convenient for her to offer something unacceptable (e.g., RP_s), then receive counteroffer RP_b, and finally re-counteroffer RP_s at time 8 (that would be surely accepted).

This anomaly arises because different types of buyers might have different optimal offers, thus revealing their private information and thus not allowing the seller to adopt the simple system of beliefs $\overline{\mu}$.

[5] Clearly, the equivalent value of a surely accepted offer of s is the offer itself.

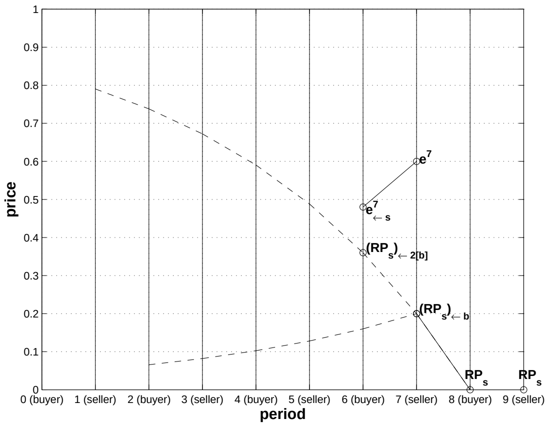

Figure 2. Non-existence of equilibrium with $RP_b = 1$, $RP_s = 0$, $\delta_b = 0.8$, $\delta_s = 0.8$, $\langle T_b = \{5, 8, 9\}, P_b^0 = \{0.5, 0.3, 0.2\}\rangle$, $T_s = 10$, $\iota(0) = b$: $x_{\leftarrow 2[b]} < e^7_{\leftarrow s}$

This happens when an "early deadline buyer" would offer the backward propagation of the seller's equivalent value at the next time point, while a "late deadline buyer" would let time pass by to reach the part of the bargaining process where backward induction prescribes her utility to be higher. Precisely, this happens when t is a possible deadline of b, the optimal strategy of s at time $t-1$ is to offer RP_b, and $U_b(x^*(t-2), t-2) < U_b(x^*(t), t)$. In other words it is not sequentially rational for b to offer $x^*(t-2)$, but it is more convenient to deviate from the $x^*(t-2)$ to offer $x^*(t)$ at t. This is the condition for the sequential rationality of strategies σ_i^*.

In Fig. 3 we depict an example of bargaining under one-sided uncertain deadlines that satisfies the above condition, so that the devised strategies are sequentially rational.

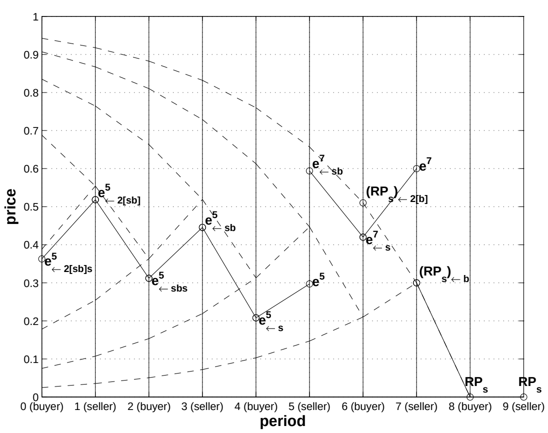

Figure 3. Backward induction with $RP_b = 1$, $RP_s = 0$, $\delta_b = 0.7$, $\delta_s = 0.7$, $\langle T_b = \{5, 8, 9\}, P_b^0 = \{0.5, 0.3, 0.2\}\rangle$, $T_s = 10$, $\iota(0) = b$: the existence is given by $e^7_{\leftarrow s} < x_{\leftarrow 2[b]}$

We can now state the following:

Theorem 3.1 *If for all t such that $\iota(t) = b$ holds $U_b(x^*(t-2), t-2) \geq U_b(x^*(t), t)$, then the assessment $a = \langle \overline{\mu}, \sigma^* \rangle$ above described is a sequential equilibrium.*

Proof sketch. The sequential rationality is easily seen from the above discussion. Consistency can be proved by the assessment sequence $a_n = \langle \mu_n, \sigma_n \rangle$ where:

- σ_n is the totally mixed strategy profile such that before the real deadline of an agent there is probability $1 - \frac{1}{n}$ of performing the action prescribed by σ and the remaining probability $\frac{1}{n}$ is uniformly distributed among the other allowed actions; while at the deadline or after it there is probability $1 - \frac{1}{n^2}$ of performing the action prescribed by σ and the remaining probability $\frac{1}{n^2}$ is uniformly distributed among the other allowed actions.
- μ_n is the system of beliefs obtained applying Bayes rule starting from the same *a priori* probability distribution P_b^0 as in $\overline{\mu}$.

Each assessment a_n is "Bayes consistent" by construction. The convergence of σ_n to σ^* is trivial. As to μ_n, given any arbitrary legal sequence S of actions (such that the bargaining does not end at the end of S and such that s is the agent acting after S), call $P_n^S = \langle \alpha^S_{n,b_1}, \alpha^S_{n,b_2}, \ldots, \alpha^S_{n,b_m} \rangle$ the probabilities that agent s assigns to b's deadlines after sequence S according to μ_n. Sequence S might contain actions that could be interpreted as being in accordance to the strategies σ^* (i.e. actions that that are the actions prescribed by strategies σ^* for some deadline T_{b_i}); be τ the time of the latest such action in S (if there are no such actions, set $\tau = -1$ by convention). Be $t = |S|$. Some calculation shows that, if $t \leq \overline{T}$, then

$$\alpha^S_{n,b_i} = \begin{cases} 0 & \text{if } T_{b_i} \leq \tau \\ \dfrac{\frac{1}{n^{t-T_{b_i}}} \cdot \alpha^0_{b_i}}{\sum\limits_{T_{b_h} \geq t} \alpha^0_{b_h} + \sum\limits_{\tau < T_{b_h} \leq t} \frac{1}{n^{t-T_{b_h}}} \cdot \alpha^0_{b_h}} & \text{if } \tau < T_{b_i} \leq t \\ \dfrac{\alpha^0_{b_i}}{\sum\limits_{T_{b_h} > t} \alpha^0_{b_h} + \sum\limits_{\tau < T_{b_h} < t} \frac{1}{n^{t-T_{b_h}}} \cdot \alpha^0_{b_h}} & \text{if } t < T_{b_i} \end{cases}$$

Therefore

$$\lim_{n \to +\infty} \alpha^S_{n,b_i} = \begin{cases} 0 & \text{if } T_{b_i} \leq t \\ \dfrac{\alpha^0_{b_i}}{\sum\limits_{T_{b_h} \geq t} \alpha^0_{b_h}} & \text{if } t < T_{b_i} \end{cases}$$

Therefore P_n^S converges to the beliefs prescribed by $\overline{\mu}$ in S. Thus μ_n converges to $\overline{\mu}$. □

4 Unsolvable Bargaining

The non-existence problem is an intrinsic limitation of the protocol. In fact, alternating-offers protocol under uncertain deadlines does not always admit pure strategy sequential equilibria. The sequential equilibrium theory assures the existence of at least one sequential equilibrium in mixed strategies [8], but the use of mixed strategy equilibria is not considered fully satisfying [13].

In this section we produce an example of bargaining that does not admit any pure strategy sequential equilibrium. The example does not even admit a *weak sequential equilibrium* [5] (a weaker notion of sequential equilibrium). Basically, an assessment is a weak sequential equilibrium if it is sequentially rational and *weakly consistent* (i.e., consistent on the equilibrium path).

Our example is depicted in Fig. 4. The b's deadline can be either at time 2 (buyer b_{early}) or at time 10 (buyer b_{late}). Shown in the figure is the backward induction construction that would have been produced by our method, but this does not lead to a solution, because if agent b's real deadline would be 10, she would not offer $e^1_{\leftarrow s}$ at

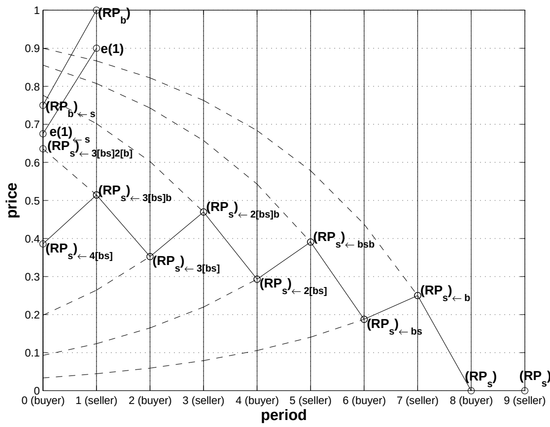

Figure 4. Non-existence of equilibrium with $RP_b = 1$, $RP_s = 0$, $\delta_b = 0.75$, $\delta_s = 0.75$, $\langle T_b = \{2, 10\}, P_b^0 = \{0.9, 0.1\}\rangle$, $T_s = 9$, $\iota(0) = b$

time $t = 0$ but would find more convenient to offer something unacceptable for s in order to offer $x_{\leftarrow 3[bs]}$ at time $t = 2$. In what follows we prove by contradiction that no pure strategy assessment works.

Theorem 4.1 *Alternating-offers bargaining with uncertain deadlines does not always admit a (weak) sequential equilibrium in pure strategies.*

Proof. Consider the bargaining of Fig. 4. By contradiction, be $a^* = \langle \sigma^*, \mu^* \rangle$ a pure strategy weak sequential equilibrium. At any time $t > 2$ the only possible system of beliefs for agent s at the equilibrium prescribes that the deadline of b is 10; therefore, for $t \geq 2$ the equilibrium strategies must be those prescribed by the backward induction construction. At time $t = 1$, instead, the beliefs of s at the equilibrium depend on which are the equilibrium actions of agent b at time $t = 0$. Two cases are possible (*i*) $\sigma^*_{b_{early}}(0) = \sigma^*_{b_{late}}(0)$ or (*ii*) $\sigma^*_{b_{early}}(0) \neq \sigma^*_{b_{late}}(0)$. We treat the two cases separately.

(*i*) If $\sigma^*_{b_{early}}(0) = \sigma^*_{b_{late}}(0)$, then by Bayes rule the beliefs of s at $t = 1$ on the equilibrium path are the initial ones (i.e., $P_b = \{0.9, 0.1\}$). Therefore, if at $t = 0$ agent b acts at the equilibrium, then at $t = 1$ the optimal strategy of agent s is to accept any offer greater than or equal to $e^1_{\leftarrow s}$ and otherwise offer RP_b. Thus, at time $t = 0$ the optimal strategy for b_{early} is to offer $e^1_{\leftarrow s}$; instead, at $t = 0$ the optimal strategy for b_{late} is to offer anything lesser than $e^1_{\leftarrow s}$ in order to be rejected and then, at $t = 2$ offer $x_{\leftarrow 3[bs]}$ that will be accepted. So $\sigma^*_{b_{early}}(0) \neq \sigma^*_{b_{late}}(0)$.

(*ii*) If $\sigma^*_{b_{early}}(0) \neq \sigma^*_{b_{late}}(0)$ then by Bayes rule the beliefs of s at $t = 1$ on the equilibrium path are: if action $\sigma^*_{b_{early}}(0)$ is observed, then $P_b = \{1, 0\}$; if action $\sigma^*_{b_{late}}(0)$ is observed, then $P_b = \{0, 1\}$. Therefore, if at $t = 0$ agent b acts at the equilibrium, then at $t = 1$ the optimal strategy of agent s is:

- if $\sigma^*_{b_{early}}(0)$ is observed, then accept any offer greater than or equal to $(RP_b)_{\leftarrow s}$ and otherwise offer RP_b;
- if $\sigma^*_{b_{late}}(0)$ is observed, then accept any offer greater than or equal to $x_{\leftarrow 4[bs]}$ and otherwise offer $x_{\leftarrow 3[bs]b}$.

Thus, at time $t = 0$ the optimal strategy for both b_{early} and b_{late} is to offer $x_{\leftarrow 4[bs]}$; so $\sigma^*_{b_{early}}(0) = \sigma^*_{b_{late}}(0)$.

In conclusion it is neither $\sigma^*_{b_{early}}(0) = \sigma^*_{b_{late}}(0)$ nor $\sigma^*_{b_{early}}(0) \neq \sigma^*_{b_{late}}(0)$. Hence the pure strategy assessment $a^* = \langle \sigma^*, \mu^* \rangle$ is not a weak sequential equilibrium. □

5 Conclusions

In this work we studied the effect of one-sided uncertain deadlines in the alternating-offers bargaining protocol with agents whose strategies are not self-constrained to any particular set of tactics. Although our study considers only one-sided uncertainty, the obtained results are particularly significant: (*1*) the equilibrium strategies can be found, when they exist, employing a simple backward procedure; (*2*) alternating-offers protocol does not always admit equilibrium in pure strategies.

The first result states that, for a wide range of bargaining parameters, the determination of equilibrium strategies of fully rational agents with one-sided uncertain deadlines can be simply tackled and the equilibrium prescribes that agreement is reached almost always at time 1.

The second result is more critical. The possible non-existence of a pure strategy equilibrium limits the possible employment of alternating-offers in real settings. Therefore the protocol should be modified (for example with the strategic delay option), or employing mixed strategies.

In future work we plan to analyze alternating-offers with strategic delay, as well as carry on the study of incomplete information alternating-offers with two-sided uncertainty over deadlines, discount factors, and reservation prices.

REFERENCES

[1] P. C. Cramton, L. M. Ausubel, and R. J. Deneckere, *Handbook of Game Theory*, volume 3, 1897–1945, Elsevier Science, 2002.

[2] F. Di Giunta and N. Gatti, 'Bargaining in-bundle multiple issues in finite-horizon alternating-offers', in *Proceedings of AIMATH*, Fort Lauderdale, USA, (January 4-6 2006).

[3] P. Faratin, C. Sierra, and N. R. Jennings, 'Negotiation decision functions for autonomous agents', *Robotic Autonomous Systems*, **24**(3-4), 159–182, (1998).

[4] S. S. Fatima, M. Wooldridge, and N. R. Jennings, 'An agenda-based framework for multi-issue negotiation', *Artificial Intelligence*, **152**, 1–45, (2004).

[5] D. Fudenberg and J. Tirole, *Game Theory*, The MIT Press, Cambridge, MA, USA, 1991.

[6] J. C. Harsanyi and R. Selten, 'A generalized Nash solution for two-person bargaining games with incomplete information', *Management Science*, **18**, 80–106, (1972).

[7] S. Kraus, *Strategic Negotiation in Multiagent Environments*, MIT Press, Cambridge, USA, 2001.

[8] D. R. Kreps and R. Wilson, 'Sequential equilibria', *Econometrica*, **50**(4), 863–894, (1982).

[9] S. Napel, *Bilateral Bargaining: Theory and Applications*, Springer-Verlag, Berlin, Germany, 2002.

[10] J. S. Rosenschein and G. Zlotkin, *Rules of Encounter. Designing Conventions for Automated Negotiations among Computers*, MIT Press, Cambridge, USA, 1994.

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

[12] A. Rubinstein, 'A bargaining model with incomplete information about time preferences', *Econometrica*, **53**(5), 1151–1172, (1985).

[13] A. Rubinstein, 'Comments on the interpretations of game theory', *Econometrica*, **59**(4), 909–924, (1991).

[14] T. Sandholm, 'Agents in electronic commerce: Component technologies for automated negotiation and coalition formation', *Autonomous Agents and Multi-Agent Systems*, **3**(1), 73–96, (2000).

[15] Tuomas W. Sandholm, 'Distributed rational decision making', in *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*, ed., Gerhard Weiss, 201–258, The MIT Press, Cambridge, MA, USA, (1999).

[16] I. Stahl, *Bargaining Theory*, Stockholm School of Economics, Stockholm, Sweden, 1972.