Bargaining with Deadlines

Tuomas Sandholm*
sandholm@cs.wustl.edu
Department of Computer Science
Washington University
One Brookings Drive
St. LouisΓMO 63130-4899

Nir Vulkan†
n.vulkan@bristol.ac.uk
Department of Economics
University of Bristol
8 Woodland Road
Bristol BS8 1TNΓEngland

Abstract

This paper analyzes automated distributive negotiation where agents have firm deadlines that are private information. The agents are allowed to make and accept offers in any order in continuous time. We show that the only sequential equilibrium outcome is one where the agents wait until the first deadline, at which point that agent concedes everything to the other. This holds for pure and mixed strategies. So, interestingly, rational agents can never agree to a nontrivial split because offers signal enough weakness of bargaining power (early deadline) so that the recipient should never accept. Similarly, the offerer knows that it offered too much if the offer gets accepted: the offerer could have done better by out-waiting the opponent. In most cases, the deadline effect completely overrides time discounting and risk aversion: an agent's payoff does not change with its discount factor or risk attitude. Several implications for the design of negotiating agents are discussed. We also present an effective protocol that implements the equilibrium outcome in dominant strategies.

1 Introduction

Multiagent systems for automated negotiation between self-interested agents are becoming increasingly important due to both technology push and application pull. For many-to-many negotiation settingsΓmarket mechanisms are often usedΓand for one-to-many negotiationΓ auctions are often appropriate. The competitive pressure on the side with many agents often reduces undesirable strategic effects. On the other handΓmarket mechanisms often have difficulty in "scaling down" to small numbers of agents (Osborne & Rubinstein 1990). In the limit of one-to-one negotiationΓstrategic considerations become prevalent. At the same timeΓone-to-one negotiation settings that crave software agents are ubiquitous. ConsiderΓfor exampleΓtwo scheduling agents negotiating meeting times on behalf of their usersΓor any e-commerce application where agents negotiate the final price of a goodΓor a scenario where agents representing different departments bargain over the details of a service which they provide jointly.

One-to-one negotiation generally involves both integrative and distributive bargaining. In integrative bargaining the agents search for *Pareto efficient* agreementsΓi.e. deals such that no other deal exists for making one agent better off without making the other worse off. IntuitivelyΓintegrative bargaining is the process of making the joint cake as large as possible. Enumerating and evaluating the Pareto efficient deals can be difficult especially in combinatorially complex settings. Automated negotiating agents hold significant promise in this arena due to their computational speed (Sandholm 1993).

In distributive bargainingΓthe focus of this paperΓthe agents negotiate on how to split the surplus provided by the dealΓi.e. how to divide the cake. A continuum of splits is possible at least if the agents can exchange sidepayments. We call any split where each agent gets a nonnegative benefit from the deal *individually rational*Γ i.e. each agent would rather accept the deal than no deal. Splitting the gains of an optimal contract in an individually rational way can be modeled generically as follows. Without loss of generalityΓthe surplus provided by the contract is normalized to 1Γand each agent's fallback payoff that would occur if no contract is made is normalized to 0. ThenΓdistributive bargaining can be studied as the process of "splitting-a-dollar". This paper focuses on designing software agents that optimally negotiate on the user's behalf in distributive bargaining.

The designer of a multiagent system can construct the interaction *protocol* (aka. mechanism) which determines the legal actions that agents can take at any point in time. Violating the protocol can sometimes be made technically impossible—e.g. disallowing a bidder from submitting multiple bids in an auction—or illegal actions can be penalized e.g. via the regular legal system. To maximize global goodΓthe protocol needs to be designed carefully taking into account that each self-interested agent will take actions so as to maximize its own utility regardless of the global good. In other wordsΓthe protocol has to provide the right incentives for the agents. In the extremeΓthe protocol could specify everythingΓi.e. give every agent at most one action to choose from at any point. HoweverΓin most negotiation settingsΓthe agents can choose whether to participate or not. SoΓto have the protocol usedΓthe designer has to provide incentives for participation as well. We will return to this question in Section 8.

The most famous model of strategic bargaining is the infinite horizon alternating offers game (Rubinstein 1982). Since it has a unique solution where agents agree on a split immediatelyΓit seems attractive for automated negotiationΓsee e.g. (KrausΓWilkenfeldΓ&

*Copyright ©, American Association of Artificial Intelligence (www.aaai.org). All rights reserved. This material is based upon work supported by the National Science Foundation under CAREER Award IRI-9703122, Grant IRI-9610122, and Grant IIS-9800994.

†Supported by the EPSRC Award GR M07052.

Zlotkin 1995). HoweverΓ the model has some weaknesses. The infinite horizon assumption is often not realistic. MoreoverΓ the results change considerably if there is a known last periodΓor even if the distribution from which the number of negotiation rounds is drawn is known and has a finite support. AlsoΓthe predictions of the model are specific to exponential time discounting and to the assumption that the dollar is infinitely divisible. For exampleΓunder linear time discounting—i.e. a fixed bargaining cost per round of offers—the results change dramatically. The first mover either gets the whole surplus or most of it depending on the ratio of the fixed costs of the two agents. The assumption of perfect information is also of limited use to designers of agents and negotiation protocols. In practiceΓagents have private information. In such settingsΓthe alternating offers model leads to multiple equilibriaΓincluding some where the true types are revealed after long delaysΓor never. The length of the delay depends on the number of types. See e.g. (Fudenberg & Tirole 1983)Γ (FudenbergΓ LevineΓ & Tirole 1985)Γ and (Rubinstein 1988). The usefulness of the model as a blueprint for designing agents and protocols is questionable when it allows for such qualitatively different outcomes.

StillΓthe tools of game theory and mechanism design can be used to study new types of bargaining modelsΓ inspired by the various applications of automated negotiation. In factΓ game theory and mechanism design theory are more suitable for software agents than for humans because agents can be designed off-line to act rationally even if determining rational strategies is complex (Rosenschein & Zlotkin 1994). AlsoΓ computational agents do not suffer from emotional irrationality. FinallyΓthe bounded rationality of computational agents can be more systematically characterized than that of humans (Sandholm & Lesser 1997).

In our modelΓagents face firm deadlines in their bargaining. This is an appealing assumption from a practical perspective since users easily understand deadlines and can trivially communicate them to software agents. Since each agent's deadline is private informationΓthere is a disadvantage in making offers. Any offer—with the exception of demanding everything for oneself—reveals some information about the proposer's deadlineΓ namely that it cannot be very long. If it wereΓ the proposer would stand a good chance of being able to out-wait the opponentΓand therefore would ask for a bigger portion of the surplus than it did. SimilarlyΓ the offerer knows that it offered too much if the offer gets accepted: the offerer could have done better by out-waiting the opponent. To model real world automated negotiationΓwe replace the temporal monopoly assumption of alternating offers at fixed intervalsΓwhich underlies Rubinstein's modelΓby a protocol where each agent can make and accept offers in continuous time. HoweverΓour results go through even if time is discrete.

Our model resembles war of attrition games where two agents compete for an object and winner-takes-all when one concedes. Those games exhibit multiplicity of equilibria. If the value of the object is common knowledgeΓthere is a symmetric equilibrium where one agent concedes immediately. There is also a symmetric equilibrium where agents concede at a rate which depends on the value of the object. (HendricksΓWeissΓ& Wilson 1988) study the general class of war of attrition games with perfect information. (Hendricks & Wilson 1989) study the case of incomplete informationΓnormally about the object's valueΓsee also (Riley 1980). The game was introduced by (Smith 1974) in a biological context and has been applied in industrial economics (Fudenberg *et al.* 1983)Γ(Kreps & Wilson 1982)).

Our model differs from the war of attrition. Agents are allowed to split the dollar instead of winner-takes-all. One would expect that to enlarge the set of equilibria and equilibrium outcomes. This intuition turns out to be false. We show that the only equilibrium outcome is one where agreement is delayed until one of the deadlines is reachedΓand then one agent gets the entire surplus.

We show that there exists a sequential equilibrium where agents do not agree to a split until the first deadlineΓat which time the agent with the later deadline receives the whole surplus. ConverselyΓ we show that there do not exist any other Bayes-Nash equilibria where agents agree to any other split at any other time. Therefore both our positive and negative results are strong with respect to the degree of sequential rationality agents are assumed to have. Intuitively speakingΓin these equilibria the agents update their beliefs rationallyΓand neither agent is motivated to change its strategy at any point of the game given that the other agent does not change its strategy.

Our results are robust in other ways as well. FirstΓ there does not exist even a mixed strategy equilibrium where an agent concedes at any rate before its deadline. ThisΓ againΓ is in contrast with the usual equilibrium analysis of war of attrition games. SecondΓthe results hold even if the agents discount time in addition to having deadlines. ThirdΓeven if agents have different risk attitudesΓ they will not agree to any split before their deadline. That isΓ even risk averse agents will refuse safe and generous offers and will instead prefer to continue to the risky "waiting game".

The rest of the paper is organized as follows. Section 2 describes our formal model of bargaining with deadlines. Section 3 presents our main results for pure-strategy equilibria. Section 4 extends them to mixed strategies. Sections 5 and 6 present the results with time discounting and with agents with risk attitudes. Section 7 describes the entailed insights for designing automated negotiating agents. Section 8 discusses implications for the design of interaction protocols. FinallyΓSection 9 concludes.

2 Our model of bargaining under deadlines

Our bargaining gameΓ $\Gamma(a, b)$Γhas two agentsΓ1 and 2. The type of agent 1 is its deadline d_1. The type of agent 2 is its deadline d_2. The types are private information: each agent only knows its own type. The type d_1 is drawn from a distribution $a(d_1)$Γand d_2 is independently drawn from a distribution $b(d_2)$Γwhere a and b are common knowledge. Without loss of generality we scale the time axis so that

$$\min(\min(d_1), \min(d_2)) = 0, \text{ and} \tag{1}$$

$$\max(\max(d_1), \max(d_2)) = 1 \tag{2}$$

That isΓ$d_1, d_2 \in (0, 1)$.

A history $\langle H(t) \rangle$ is a (possibly empty) list of all unaccepted offers up to time t. An *action* at time t describes the threshold of what offers are acceptable to the agent from then onΓand what offer the agent has outstanding from then on. We do not assume that the agents have to improve their offers over time. They can even offer less than they offered earlier. Each agent also has a belief updating rule. FinallyΓeach agent chooses a strategy for the game. A *pure strategy* is an agent's deterministic mapping from her history to actionΓi.e. it defines what offers the agent would accept and make as a function of what offers and rejections/acceptances the other agent has made so far. A *mixed strategy* is an agent's probability mixture over pure strategiesΓi.e. the agent can secretly throw a (possibly biased) dice to choose up front which pure strategy to use.

The agents can agree to any split $(x, 1-x)$ where $x \in [0, 1]$. The payoffs from an agreement $(x, 1-x)$ at time $t < d_1, d_2$ are x and $1-x$ for agents 1 and 2 respectively.[1] The payoff for an agent from any agreement which takes place after her deadline is 0. We assume that an agent strictly prefers to hand over the whole surplus than to miss her deadline. In other wordsΓ if the agent will get zero payoff anywayΓit will rather give the surplus to the other agent than not.[2]

We use two standard game-theoretic solution concepts to model how rational agents would play the game: Bayes-Nash equilibrium and sequential equilibrium. In briefΓa strategy-profile (s_1, s_2) is a *Bayes-Nash equilibrium* if s_1 is best response to s_2 in every information setΓs_2 is best response to s_1 in every information setΓand each agent updates her beliefs based on the strategies and the observed history using Bayes rule (Mas-ColellΓWhinstonΓ& Green 1995).

A *sequential equilibrium* is a refinement of Bayes-Nash equilibrium which places the further requirement that agents act rationally also in zero probability information sets. Intuitively speakingΓagents must not use threats that are not credible. FormallyΓa strategy-profile (s_1, s_2) is a sequential equilibrium if (s_1, s_2) is a Bayes-Nash equilibrium and there exists a sequence (s_1^n, s_2^n) such that (i) for any nΓs_1^n and s_2^n consist of beliefs which are fully mixedΓi.e. beliefs which attach positive probability to every information setΓand (ii) (s_1^n, s_2^n) converges to (s_1, s_2) (Mas-ColellΓWhinstonΓ& Green 1995).

3 Pure strategy equilibria

It turns out that a "sit-and-wait" way of playing the game is one rational way for agents to behave:

Proposition 3.1 *There exists a sequential equilibrium of* $\Gamma(a, b)$*, where the agent with the latest deadline receives the whole surplus exactly at the earlier deadline.*

Proof. Consider the following (symmetric) strategies: $s_i(d_i)$: demand $x = 1$ (everything) at any time $t < d_i$. At $t = d_i$ accept any offer. At any time $t < d_i$Γ reject offers $x_i \geq d_i$. Update beliefs according to Bayes ruleΓputting all the weight of the posterior distribution over the values of the opponent's deadlineΓd_{-i}Γover the interval $(1 - x_i, 1)$. At any time $t < d_i$Γreject offers $x_i < d_i$. Update beliefs in the following way. If $1 - x_i$ is already in the prior of d_{-i} then the posterior is simply equal to the prior. If $1 - x_i$ is not in the prior of d_{-i} then add the point $1 - x_i$ and update according to Bayes rule.

We first show that the beliefs specified above are consistent: Let s_i^ϵ (for $i = 1, 2$) assign probability $1 - \epsilon$ to the above specified posterior beliefsΓand probability ϵ to the rest of the support of d_{-i}. As $\epsilon \to 0$ the fully mixed strategy pair converges to (s_1, s_2)Γwhile the beliefs generated by the fully mixed strategy pair converge to the beliefs described above. It is now easy to see thatΓgiven these beliefsΓactions are sequentially rational. Along the equilibrium path agents will always demand the full surplus and therefore no agreement will be achieved before the fist deadline. At the first deadlineΓthe agent with the later deadline will receive the whole surplus. □

The equilibrium described above is clearly also an equilibrium of any subgame beginning at time t onwards regardless of the historyΓ$\langle H(t) \rangle$. That isΓat any stage of the game each agent can move to this waiting gameΓ and the other agent's best response is to do the same. We use this property in the proof of our main result which states thatΓsurprisinglyΓthe "sit-and-wait" way of playing the game is the only rational one. We could prove this impossibility result using sequential equilibrium. HoweverΓwe prove the stronger claim that there is no equilibrium even using a weaker definition of equilibriumΓthe Bayes-Nash equilibrium. It follows that no sequential equilibrium exists either. For pedagogical reasons we present the result for pure strategies first. We generalize it to mixed strategies in Section 4.

Theorem 3.1 *If* $d_1 > 0$ *or* $d_2 > 0$*, there does not exist a pure strategy Bayes-Nash equilibrium of* $\Gamma(a, b)$*, where agents agree to a split other than* $(1, 0)$ *or* $(0, 1)$.

[1] Payoffs could be defined in several ways. One approach, by Smith (1997), is to define payoffs in terms of aspiration levels, or acceptance thresholds. However, it turns out to be sufficient for the results of this paper to express payoffs not in terms of strategies, but in the traditional way of stating them in terms of preferences over final agreements.

[2] Our results still hold if this assumption is removed, but an additional set of equilibria appears where agents miss their deadline with probability 1. Still, our assumption can be defended on the grounds that there is usually a small (possibly infinitesimally small), but positive utility associated with reaching an agreement.

Proof. AssumeΓfor contradictionΓthat there exist types $d_1 > 0$ and $d_2 > 0$ and a pure strategy Bayes-Nash equilibrium (s_1, s_2) where the agents agree to a split $(\pi_1, \pi_2) = (x, 1-x)$Γwhere $x \in (0,1)$ at time $t \geq 0$. We assumeΓwithout loss of generalityΓthat agent 1 receives at least one halfΓi.e. $x \geq \frac{1}{2}$. We can therefore write $x = \frac{1}{2} + \epsilon$Γwhere $\frac{1}{2} > \epsilon \geq 0$.

Let $g_0(d_2)$ denote agent 1's beliefs about d_2 at time t. SimilarlyΓlet $f_0(d_1)$ denote agent 2's beliefs about d_1 at time t. Denote by $G(d_2) \equiv \int_t^1 g_0(d_2)dd_2$ the cumulative distribution of g_0Γand by $F(d_1) \equiv \int_t^1 f_0(d_1)dd_1$ the cumulative distribution of f_0.

In equilibriumΓagent 2 will accept $1-x$ only if she does not expect to receive more by unilaterally moving to the waiting game. The expected payoff from the waiting game is simply agent 2's subjective probability that $d_2 < d_1$. Hence agent 2 would accept only if

$$\frac{1}{2} - \epsilon \geq \int_t^{d_2} f_0(d_1)dd_1 = F(d_2) - F(t) = F(d_2) \quad (3)$$

In other wordsΓagent 2's typeΓd_2Γmust not be too high. Let $\alpha(y) \equiv \inf[d_2 \mid y \geq F(d_2)]$. With this notationΓ(3) can be rewritten as $d_2 \leq \alpha(\frac{1}{2} - \epsilon)$.

NowΓagent 1 will only accept this offer if it will give her an expected payoff at least as large as that of the waiting gameΓwhich equals her subjective probability of winning the waiting game. There are two cases. FirstΓif $d_1 > \alpha(\frac{1}{2} - \epsilon)$Γagent 1 knows that she will win the waiting game with probability oneΓso the split $(x, 1-x)$ could not occur in equilibrium. The second case occurs when $d_1 \leq \alpha(\frac{1}{2} - \epsilon)$. Agent 1 can use the fact that $d_2 \leq \alpha(\frac{1}{2} - \epsilon)$ to update her beliefs about agent 2's deadline as follows:

$$\begin{aligned} g_1(d_2) &= 0 \text{ if } d_2 > \alpha(\frac{1}{2} - \epsilon). \text{ Otherwise}\Gamma \\ g_1(d_2) &= g_0(d_2)\frac{\int_t^1 g_0(d_2)}{\int_t^{\alpha(\frac{1}{2}-\epsilon)} g_0(d_2)} \\ &= g_0(d_2) \underbrace{[1 + \frac{\int_{\alpha(\frac{1}{2}-\epsilon)}^1 g_0(d_2)}{\int_t^{\alpha(\frac{1}{2}-\epsilon)} g_0(d_2))}]}_{\geq 2 \text{ because } \alpha(\frac{1}{2}-\epsilon) \leq \text{median}(g_0)} \end{aligned} \quad (4)$$

Based on these updated beliefsΓagent 1 would accept only if

$$\frac{1}{2} + \epsilon \geq \int_t^{d_1} g_1(d_2)dd_2 \geq 2[G(d_1) - G(t)] = 2G(d_1) \quad (5)$$

In other wordsΓagent 1's typeΓd_1Γmust also not be too high. Let $\beta(y) \equiv \inf[d_1 \mid y \geq G(d_1)]$. With this notationΓ(5) can be rewritten as $d_1 \leq \beta(\frac{\frac{1}{2}+\epsilon}{2})$.

In equilibriumΓagent 2 only accepts if it gives her an expected payoff at least as large as that of the waiting gameΓwhich equals her subjective probability of winning the waiting game. There are two cases. FirstΓif $d_2 > \beta(\frac{\frac{1}{2}+\epsilon}{2})$Γagent 2 knows that she will win the waiting game with probability oneΓso the split $(x, 1-x)$ could not occur. The second case occurs when $d_2 \leq \beta(\frac{\frac{1}{2}+\epsilon}{2})$. Agent 2 can use the fact that $d_1 \leq \beta(\frac{\frac{1}{2}+\epsilon}{2})$ to update her beliefs about agent 1's deadline as follows:

$$\begin{aligned} f_1(d_1) &= 0 \text{ if } d_1 > \beta(\frac{\frac{1}{2}+\epsilon}{2}). \text{ Otherwise}\Gamma \\ f_1(d_1) &= f_0(d_1)\frac{\int_t^1 f_0(d_1)}{\int_t^{\beta(\frac{\frac{1}{2}+\epsilon}{2})} f_0(d_1)} \\ &= f_0(d_1) \underbrace{[1 + \frac{\int_{\beta(\frac{\frac{1}{2}+\epsilon}{2})}^1 f_0(d_1)}{\int_t^{\beta(\frac{\frac{1}{2}+\epsilon}{2})} f_0(d_1))}]}_{\geq 2 \text{ because } \beta(\frac{\frac{1}{2}+\epsilon}{2}) \leq \text{median}(f_0)} \end{aligned} \quad (6)$$

Based on these updated beliefsΓagent 2 would accept only if

$$\frac{1}{2} - \epsilon \geq \int_t^{d_2} f_1(d_1)dd_1 \geq 2[F(d_2) - F(t)] = 2F(d_2) \quad (7)$$

i.e. $d_2 \leq \alpha(\frac{\frac{1}{2}-\epsilon}{2}))$.

This process of belief update and acceptance threshold resetting continues to alternate between agents. After r rounds of this alternationΓall types have been eliminated except those that satisfy $d_1 \leq \beta(\frac{\frac{1}{2}-\epsilon}{2^r})$ and $d_2 \leq \alpha(\frac{\frac{1}{2}+\epsilon}{2^r})$. This process can continue for an unlimited number of stepsΓrΓso the upper bounds approach zero. Therefore the equilibrium cannot exist if $d_1 > 0$ or $d_2 > 0$. Contradiction. □

4 Mixed strategy equilibria

We now strengthen our impossibility result by showing that it holds for mixed strategies as wellΓi.e. that there is no other rational way of playing the game than "sit-and-wait" even if randomization is possible. This is yet another difference between our setting and war of attrition games. In the latterΓmixed strategies play an important role: typically the unique symmetric equilibrium has concession rates that are mixed strategies.

Theorem 4.1 *If $d_1 > 0$ or $d_2 > 0$, there does not exist a mixed strategy Bayes-Nash equilibrium of $\Gamma(a, b)$, where the agents agree to a split other than $(1,0)$ or $(0,1)$ with positive probability.*

Proof. AssumeΓfor contradictionΓthat there exist types $d_1 > 0$ and $d_2 > 0$ and a mixed strategy Bayes-Nash equilibrium where there is positive probability of an agreement other than $(1,0)$ or $(0,1)$. Now there has to exist at least one point in timeΓtΓwhere there is positive probability of an agreement other than $(1,0)$ or $(0,1)$. We analyze the equilibrium at such a time t. Recall fΓgΓFΓGΓαΓand β from the proof of Thrm. 3.1.

Agent 1 will accept an agreement if she gets a share $x \geq a_1$Γwhere a_1 is her acceptance threshold. That threshold depends on her type. Since we are analyzing a mixed strategy equilibriumΓthe threshold can also depend on randomization. We therefore say that a_1 is randomly chosen for time t from a probability density function $m(a_1)$. SimilarlyΓagent 2 will accept an agreement if she has to offer 1 a share of $x \leq a_2$ where a_2 is agent 2's offering threshold. We say that a_2 is chosen for time t from a probability density function $n(a_2)$.

Without loss of generalityΓwe assume that there is positive probability that the agreement is made in the range $x \geq \frac{1}{2}$. This implies that there is positive probability that $a_2 \geq \frac{1}{2}$.

Let $\underline{a_1}$ be the smallest a_1 in the support of m (alternatively let $\underline{a_1}$ be the infimum of m). The assumption that there is positive probability of an agreement other than $(1,0)$ or $(0,1)$ means that $\underline{a_1} = 1 - \epsilon$ for some $\epsilon > 0$.

Because the strategies are in equilibriumΓm and n must be best responses to each other. For n to be a best responseΓeach thresholdΓa_2Γin the support of n has to give agent 2 at least the same payoff as she would get by going to the waiting game. Focusing on those a_2 for which $a_2 \geq \frac{1}{2}$ this means

$$\frac{1}{2} \geq E[\pi_2^{wait}] = \int_t^{d_2} f_0(d_1)dd_1 = F(d_2) - F(t) = F(d_2) \tag{8}$$

SoΓ$d_2 \leq \alpha(\frac{1}{2})$.

NowΓin equilibriumΓevery strategy in the support of m has to give agent 1 at least the same payoff that she would get by going to the waiting game. There are two cases. FirstΓif $d_1 > \alpha(\frac{1}{2})$Γagent 1 knows that she will win the waiting game with probability oneΓso the split $(x, 1-x)$ could not occur in equilibrium. The second case occurs when $d_1 \leq \alpha(\frac{1}{2})$. Agent 1 can use the fact that $d_2 \leq \alpha(\frac{1}{2})$ to update her beliefs about agent 2's deadline:

$$\begin{aligned} g_1(d_2) &= 0 \text{ if } d_2 > \alpha(\frac{1}{2}). \text{ Otherwise}\Gamma \\ g_1(d_2) &= g_0(d_2)\frac{\int_t^1 g_0(d_2)}{\int_t^{\alpha(\frac{1}{2})} g_0(d_2)} \\ &= g_0(d_2) \underbrace{[1 + \frac{\int_{\alpha(\frac{1}{2})}^1 g_0(d_2)}{\int_t^{\alpha(\frac{1}{2})} g_0(d_2))}]}_{\geq 2 \text{ because } \alpha(\frac{1}{2}) \leq \text{ median}(g_0)} \end{aligned} \tag{9}$$

Based on these updated beliefsΓthe support of m can include $\underline{a_1}$ only if

$$\underline{a_1} = 1 - \epsilon \geq \int_t^{d_1} g_1(d_2)dd_2 \geq 2[G(d_1) - G(t)] = 2G(d_1) \tag{10}$$

In other wordsΓagent 1's typeΓd_1Γcannot be too high. SpecificallyΓthis can be written as $d_1 \leq \beta(\frac{1-\epsilon}{2})$.

In equilibriumΓevery strategy in the support of n has to give agent 2 at least the same payoff that she would get by going to the waiting gameΓwhich equals her subjective probability of winning the waiting game. There are two cases. FirstΓif $d_2 > \beta(\frac{1-\epsilon}{2})$Γagent 2 knows that she will win the waiting game with probability oneΓso the split $(x, 1-x)$ could not occur. The second case occurs when $d_2 \leq \beta(\frac{1-\epsilon}{2})$. Agent 2 can use the fact that $d_1 \leq \beta(\frac{1-\epsilon}{2})$ to update her beliefs about agent 1's deadline:

$$\begin{aligned} f_1(d_1) &= 0 \text{ if } d_1 > \beta(\frac{1-\epsilon}{2}). \text{ Otherwise}\Gamma \\ f_1(d_1) &= f_0(d_1)\frac{\int_t^1 f_0(d_1)}{\int_t^{\beta(\frac{1-\epsilon}{2})} f_0(d_1)} \\ &= f_0(d_1) \underbrace{[1 + \frac{\int_{\beta(\frac{1-\epsilon}{2})}^1 f_0(d_1)}{\int_t^{\beta(\frac{1-\epsilon}{2})} f_0(d_1))}]}_{\geq 2 \text{ because } \beta(\frac{1-\epsilon}{2}) \leq \text{ median}(f_0)} \end{aligned} \tag{11}$$

Based on these updated beliefsΓand focusing on those a_2 for which $a_2 \geq \frac{1}{2}$ we can rule out high values of d_2 (otherwise agent 2 would be better off by waiting):

$$\frac{1}{2} \geq E[\pi_2^{wait}] = \int_t^{d_2} f_1(d_1)dd_1 \geq 2F(d_2) \tag{12}$$

i.e. $d_2 \leq \alpha(\frac{1}{4})$).

This process of belief update and acceptance threshold resetting continues to alternate between agents. After r rounds of this alternationΓall types have been eliminated except those that satisfy $d_1 \leq \beta(\frac{1-\epsilon}{2^r+1})$ and $d_2 \leq \alpha(\frac{1}{2^r+1})$. This process can continue for an unlimited number of stepsΓrΓso the upper bounds approach zero. Therefore the equilibrium cannot exist if $d_1 > 0$ or $d_2 > 0$. Contradiction. □

5 Incorporating discounting

Time discounting is a standard way of modeling settings where the value of the good decreases over timeΓe.g. due to inflation or due to perishing. In the previous sections we assumed that agents do not discount time. HoweverΓwe now show that our results are robust to the case where agents do discount time in addition to having firm deadlines. Let δ_1 be the discount factor of agent 1Γand δ_2 be the discount factor of agent 2. The utility of agent i from an agreement where he receives a share x at time $t < d_i$ is then $\delta_i^t x$. We denote by $\Gamma(a, b, \delta_1, \delta_2)$ the bargaining game where aΓbΓδ_1Γand δ_2 are common knowledge. We now prove that our previous result for $\Gamma(a, b)$ holds also for a large range of parameters in $\Gamma(a, b, \delta_1, \delta_2)$. SoΓinterestinglyΓthe bargaining power of an agent does not change with her discount factorΓin contrast to the results of most other bargaining games. In other wordsΓthe deadline effect completely suppresses the discounting effect. This crisp result is important in its own right for the design of automated negotiating agentsΓand it also motivates the study of deadline-based models as opposed to focusing only on discounting-based ones.

Proposition 5.1 *For any $\delta_1, \delta_2, 0 < \delta_1 \leq 1, 0 < \delta_2 \leq 1$, there exists a sequential equilibrium of $\Gamma(a, b, \delta_1, \delta_2)$ where the agent with the latest deadline receives the whole surplus exactly at the earlier deadline.*

Proof. The equilibrium strategies and proof of sequential equilibrium are identical to those in the proof of Proposition 3.1 with the difference that the threshold is no longer d_i but $\delta_i^t d_i$. Also the posteriors are now defined only until δ_i^t and not 1. □

Theorem 5.1 *If $\delta_1\delta_2 > \frac{1}{2}$, there does not exist a Bayes-Nash equilibrium of $\Gamma(a, b, \delta_1, \delta_2)$, (in pure or mixed strategies) where agents agree to a split other than $(1, 0)$ or $(0, 1)$.*

Proof. We prove the case for pure strategy equilibrium. The extension for mixed strategy equilibrium is identical to that in Theorem 4.1. The proof is a variant of the proof in Theorem 3.1Γand we keep the notation from there.

AssumeΓ for contradictionΓ that there exist types $d_1 > 0$ and $d_2 > 0$ and a pure strategy Bayes-Nash equilibrium (s_1, s_2) where the agents agree to a split of the total surplus available at time tΓaccording to proportions $(\pi_1, \pi_2) = (x, 1-x)$Γwhere $x \in (0, 1)$ at time $t \geq 0$. We assumeΓwithout loss of generalityΓthat agent 1 receives at least one halfΓi.e. $x \geq \frac{1}{2}$. We can therefore write $x = \frac{1}{2} + \epsilon$Γwhere $\frac{1}{2} > \epsilon \geq 0$.

In equilibriumΓagent 2 will accept $1 - x$ only if she does not expect to receive more by unilaterally moving to the waiting game. The expected payoff from the waiting game is now agent 2's subjective probability that $d_2 < d_1$Γ multiplied by the discounted value of winning. Hence agent 2 would accept only if

$$\delta_2^t(\frac{1}{2}-\epsilon) \geq \int_t^{d_2} \delta_2^{d_1} f_0(d_1)dd_1 \geq \delta_2 \int_t^{d_2} f_0(d_1)dd_1 \quad (13)$$

Dividing both sides by δ_2^tΓwe get:

$$\frac{1}{2} - \epsilon \geq \delta_2^{1-t} \int_t^{d_2} f_0(d_1)dd_1 = \delta_2^{1-t}F(d_2) \geq \delta_2 F(d_2) \quad (14)$$

In other words: $d_2 \leq \alpha\frac{(\frac{1}{2}-\epsilon)}{\delta_2}$.

NowΓagent 1 can use this to update her beliefs about agent 2's deadline in the same way as in equation (4)Γ with the difference that now $g_1 \geq 2\delta_2 g_0$. Since $\delta_2 > 0.5$ (by the assumption that $\delta_1\delta_2 > 0.5$)Γ we know that $g_1 > g_0$ (when g_1 is not zero).

Based on these updated beliefsΓagent 1 would accept only if

$$\delta_1^t(\frac{1}{2}+\epsilon) \geq \int_t^{d_1} \delta_1^{d_2} g_1(d_2)dd_2 \geq \delta_1 \int_t^{d_1} g_1(d_2)dd_2 \quad (15)$$

Dividing both sides by δ_1^t and using the updated beliefsΓg_1Γwe can now rule out "high" types of agent 1. FormallyΓ$d_1 \leq \beta(\frac{\frac{1}{2}+\epsilon}{2\delta_2\delta_1})$.

Once moreΓbelief updating by agent 2 (in the same way as in (6)) yields $f_1 \geq 2\delta_2\delta_1 f_0$. Since $\delta_1\delta_2 > 0.5$ we get that $f_1 > f_0$ (when it is not zero). Based on these updated beliefsΓagent 2 would accept only if

$$\delta_2^t(\frac{1}{2}-\epsilon) \geq \int_t^{d_2} \delta_2^{d_1} f_1(d_1)dd_1 \geq \delta_2 \int_t^{d_2} f_0(d_1)dd_1 \quad (16)$$

Dividing by δ_2^t and using the updated beliefsΓf_1 we can rule out the following types: $d_2 \leq \alpha(\frac{\frac{1}{2}-\epsilon}{2\delta_2^2\delta_1}) < \alpha(\frac{\frac{1}{2}-\epsilon}{2\delta_2})$. This process of belief update and acceptance threshold resetting continues to alternate between agents. After r rounds of this alternationΓ all types have been eliminated except those that satisfy $d_1 \leq \beta(\frac{\frac{1}{2}-\epsilon}{(2\delta_1)^r})$ and $d_2 \leq \alpha(\frac{\frac{1}{2}+\epsilon}{(2\delta_2)^r})$. This process can continue for an unlimited number of stepsΓrΓso the upper bounds approach zero. Therefore the equilibrium cannot exist if $d_1 > 0$ or $d_2 > 0$. Contradiction. □

For exampleΓif the annual interest rate is 10%Γthe discount factor would be $\delta = \frac{1}{1+0.1} \approx 0.909$ per year. For the conditions of Theorem 5.1 to be violatedΓat least one agent's discount factor would have to be $\delta_i \leq \frac{1}{2}$. This would mean that the unit of time from which its deadline is drawn would have to be no shorter than 7 years because $\frac{1}{2} < (\frac{1}{1+0.1})^7$. Since most deadline bargaining situations will certainly have shorter deadlines than 7 yearsΓTheorem 5.1 shows that "sit-and-wait" is the only rational strategy. SoΓin practiceΓthe effect of deadlines suppresses that of discount factors. This is even more commonly true in automated negotiation because that is most likely going to be used mainly for fast negotiation at the operative decision making level instead of strategic long-term negotiation.

6 Robustness to risk attitudes

We now generalize our results to agents that are not necessarily risk neutral. Usually in bargaining games the equilibrium split of the surplus depends on the agents' risk attitudes. HoweverΓwe show that this does not happen in our setting. This is surprising at first since a risk averse agent generally prefers a smaller but safe share to the risky option of the waiting game even if she expects to win it with high probability. HoweverΓ we show that the type-elimination effect described in the theorems so far is still present and dominates any concessions which may be consistent with risk aversion.

Let the agents' risk attitudes be captured by utility functionsΓu_i where $i = 1, 2$. Without loss of generality we let $u_i(0) = 0$ and $u_i(1) = 1$ for both agents.

Proposition 6.1 *There exists a sequential equilibrium of $\Gamma(a, b, u_1, u_2)$, where the agent with the latest deadline receives the whole surplus exactly at the earlier deadline.*

Proof. The equilibrium strategies and proof of sequential equilibrium are identical to those in the proof of Proposition 3.1 with the difference that the threshold is no longer d_i but $u_i(d_i)$. □

The following definition is used to state our main result for the case of different risk attitudes.

Definition. 6.1 *The* maximum risk aversion *of agent i is*

$$\rho_i \equiv \max_x \frac{u_i(x)}{x} \tag{17}$$

We can now show that our impossibility result applies to a large range of risk attitudes of the agents:

Theorem 6.1 *If $\rho_1\rho_2 < 2$, there does not exist a Bayes-Nash equilibrium (pure or mixed) of $\Gamma(a, b, u_1, u_2)$, where agents agree to a split other than $(1, 0)$ or $(0, 1)$.*

Proof. We prove the case for pure strategy equilibrium. The extension to mixed strategies is identical to that in Theorem 4.1. We keep the notations from Theorem 3.1. AssumeΓfor contradictionΓthat there exist types $d_1 > 0$ and $d_2 > 0$ and a pure strategy Bayes-Nash equilibrium (s_1, s_2) where the agents agree to a split $(\pi_1, \pi_2) = (x, 1-x)$Γwhere $x \in (0, 1)$ at time $t \geq 0$. AssumeΓwithout loss of generalityΓthat agent 1 receives at least halfΓi.e. $x \geq \frac{1}{2}$. Thus we can write $x = \frac{1}{2} + \epsilon$Γ where $\frac{1}{2} > \epsilon \geq 0$. In equilibriumΓagent 2 accepts $1 - x$ only if she does not expect to receive more by unilaterally moving to the waiting game. The expected payoff from the waiting game is agent 2's subjective probability that $d_2 < d_1$. SoΓagent 2 would accept only if

$$\rho_2(\frac{1}{2} - \epsilon) \geq u_2(\frac{1}{2} - \epsilon) \geq \int_t^{d_2} f_0(d_1)dd_1 = F(d_2) \tag{18}$$

In other wordsΓ$d_2 \leq \alpha(\rho_2(\frac{1}{2} - \epsilon))$. NowΓagent 1 can use this to update her beliefs about agent 2's deadline in the same way as in equation (4)Γwith the difference that now $g_1 \geq \frac{2}{\rho_2} g_0$. Since (by assumption) $\rho_2 < 2$Γ then $g_1 > g_0$ (when g_1 is not zero). Based on these updated beliefsΓagent 1 would accept only if:

$$u_1(\frac{1}{2} + \epsilon) \geq \int_t^{d_1} g_1(d_2)dd_2 \tag{19}$$

Using ρ_1 and the updated beliefsΓg_1Γwe can rule out "high" types of agent 1. FormallyΓ$d_1 \leq \beta(\rho_2\rho_1(\frac{1}{2} + \epsilon))$.

Once moreΓbelief updating by agent 2 (in the same way as in (6)) yields $f_1 \geq \frac{2}{\rho_2\rho_1} f_0$. Since $\rho_2\rho_1 < 2$ we get $f_1 > f_0$ (when f_1 is not zero). Based on these updated beliefsΓagent 2 would accept only if

$$u_2(\frac{1}{2} - \epsilon) \geq \int_t^{d_2} f_1(d_1)dd_1 \tag{20}$$

Using ρ_2 and the updated beliefsΓf_1Γwe can rule out the following types: $d_2 \leq \alpha(\rho_2^2\rho_1(\frac{1}{2} - \epsilon)) < \alpha(\rho_2(\frac{1}{2} - \epsilon))$. This process of belief update and acceptance threshold resetting continues to alternate between agents. After r rounds of this alternationΓall types have been eliminated except those that satisfy $d_1 \leq \beta(\rho_1^r(\frac{1}{2} - \epsilon))$ and $d_2 \leq \alpha(\rho_2^r(\frac{1}{2} - \epsilon))$. This process can continue for an unlimited number of stepsΓrΓso the upper bounds approach zero. Therefore the equilibrium cannot exist if $d_1 > 0$ or $d_2 > 0$. Contradiction. □

7 Designing bargaining agents

Our motivation for studying bargaining with deadlines stems from our desire to construct software agents that will optimally negotiate on behalf of the real world parties that they represent. That will put experienced and poor human negotiators on an equal footingΓand save human negotiation effort.

Deadlines are widely advocated and used in automated electronic commerce to capture time preference. For example when a user delegates priceline.com to find an inexpensive airline flight on the webΓthe user gives it one hour to complete (while priceline.com uses an agent with a deadlineΓthe setting is a form of auctionΓ not bargaining). Users easily understand deadlinesΓand it is simple to specify a deadline to an agent.

Our results show that in distributive bargaining settings with two agents with deadlinesΓit is not rational for either agent to make or accept offers. But what if a rational software agent receives an offer from the other party? This means that the other party is irrationalΓand could perhaps be exploited. HoweverΓthe type-elimination argument from the proofs above applies here tooΓand it is not rational for the software agent to accept the offerΓno matter how good it is. To exploit the other partyΓthe agent would have to have an opponent model to model the other party's irrationality. While game theory allows us to give precise prescriptions for rational playΓit is mostly silent about irrationalityΓand how to exploit it.

Another classic motivation for automated negotiation is that computerized agents can negotiate faster. HoweverΓin distributive bargaining settings where the agents have deadlinesΓthis argument does not hold because in such settingsΓrational software agents would sit-and-wait until one of the deadlines is reached. From an implementation perspective this suggests the use of daemons that trigger right before the deadline instead of agents that use computation before the deadline.

FinallyΓour results suggest that a user will be in a much stronger bargaining position by inputting time preferences to her agent in terms of a time discount function instead of a deadlineΓeven if the discounting is significant. To facilitate thisΓsoftware agent vendors should provide user interfaces to their agents that allow easy human-understandable specification of time discounting functions instead of inputting a deadline.

8 Designing bargaining protocols

The following mechanism implementsΓin dominant strategiesΓthe equilibrium of the deadline bargaining game described above. FirstΓagents report their deadlinesΓ$\tilde{d}_i$Γto the protocol—possibly insincerely ($\tilde{d}_i \neq d_i$). The protocol then assigns the whole dollar to the agent with the highest $\tilde{d}_i$Γbut this only takes place at time $t = \tilde{d}_{-i}$Γi.e. at the earlier reported deadline. Truth-telling is a dominant strategy in this mechanism. By reporting $\tilde{d}_i < d_i$ agent i's probability of winning is reduced. By reporting $\tilde{d}_i > d_i$Γagent i increase its probability of winningΓbut only in cases where $\tilde{d}_{-i} > d_i$Γi.e. when i misses its deadline. ThereforeΓreporting $\tilde{d}_i = d_i$ is a dominant strategy.

This mechanism is efficient in several ways. FirstΓ it minimizes counterspeculation. In equilibria that are based on refinements of the Nash equilibrium—such as Bayes-Nash or sequential equilibrium—an agent's best strategy depends on what others will do. This requires speculation about the others' strategiesΓwhich can be speculated by considering the others' typesΓtheir rationalityΓwhat they think of the former agentΓwhat they think the former agent thinks about themΓ*etc. ad infinitum.* On the other handΓin a dominant strategy mechanism an agent's strategy is optimal no matter what others do. ThereforeΓcounterspeculation is not useful. The agent need not waste time in counterspeculation which can be intractable or even noncomputable. In additionΓit is easier to program an agent that executes a given dominant strategy than an agent that counterspeculates. SecondΓdominant strategy mechanisms are robust against irrational agents since their actions do not affect how others should behave. FinallyΓ the mechanism minimizes communication: each agent only sends one message.

HoweverΓthe mechanism is not Pareto efficient if time is being discountedΓbecause the agreement is delayed as it was in the original free-form bargaining game. In such settingsΓany mechanism that results in an immediate agreement is Pareto efficientΓe.g. a protocol that forces a 50:50 split up front. This protocol is efficient in all respects discussed above. It might seem like a good solution to the problem raised by our impossibility results. HoweverΓagents in e-commerce applications usually can choose whether to use a protocol or not. If agents know their types before they choose the protocol they want to useΓan adverse selection problem arises. To see whyΓassume that types are normally distributed. This is assumed for simplicity of presentation and is not crucial. Agents with deadlines above 1/2 will not participate in such a protocol because they can expect to do better in a free-form bargaining setting. But if only agents with deadlines below 1/2 participateΓagents with deadlines between 1/4 and 1/2 should not participate. NextΓagents with deadlines between 1/8 and 1/4 would not participateΓand so on. In equilibriumΓno agent would participate. This argument does not rely on a 50:50 split. The adverse selection problem will affect any protocol that does not implement Nash (or stronger) equilibrium outcomes.

9 Conclusions

Automated agents have been suggested as a way to facilitate increasingly efficient negotiation. In settings where the bargaining setΓi.e. set of individually rational Pareto efficient dealsΓis difficult to construct for example due to a combinatorial number of possible deals (Sandholm 1993) or the computational complexity of evaluating any given deal (Sandholm & Lesser 1997)Γ the computational speed of automated agents can significantly enhance negotiation. Additional efficiency can stem from the fact that computational agents can negotiate with large numbers of other agents quickly and virtually with no negotiation overhead. HoweverΓ this paper showed that in one-to-one negotiation where the optimal deal in the bargaining set has been identified and evaluatedΓand distributing the profits is the issueΓan agent's power does not stem from speedΓbut on the contraryΓfrom the ability to wait.

We showed that in one-to-one bargaining with deadlinesΓthe only sequential equilibrium is one where the agents wait until the first deadline is reached. This is in line with some human experiments where adding deadlines introduced significant delays in reaching agreement (RothΓMurnighanΓ& Schoumaker 1988). We also showed that deadline effects almost always completely suppress time discounting effects. Impossibility of an interim agreement also applies to most types of risk attitudes of the agents. The results show that for deadline bargaining settings it is trivial to design the optimal agent: it should simply wait until it reaches its deadline or the other party concedes. On the other handΓ a user is better off by giving her agent a time discount function instead of a deadline since a deadline puts her agent in a weak bargaining position. FinallyΓwe discussed mechanism designΓand presented an effective protocol that implements the outcome of the free-form bargaining game in dominant strategy equilibrium.

References

Fudenberg, D., and Tirole, J. 1983. Sequential bargaining with incomplete information. *Review of Economic Studies* 50:221–247.

Fudenberg, D.; Glibert, R.; Stiglitz, J.; and Tirole, J. 1983. Preemption, leapfrogging, and competition in patent races. *European Economic Review* 22:3–31.

Fudenberg, D.; Levine, D.; and Tirole, J. 1985. Infinite-horizon models of bargaining with one-sided incomplete information. In Roth, ed., *Game-Theoretic Models of Bargaining.* Cambridge U. Press. 73–98.

Hendricks, K., and Wilson, C. 1989. The war of attrition in discrete time. Technical report, State University of New York, Stony Brook. Mimeo.

Hendricks, K.; Weiss, A.; and Wilson, C. 1988. The war of attrition in continuous time with complete information. *Internat. Economic Review* 29:663–680.

Kraus, S.; Wilkenfeld, J.; and Zlotkin, G. 1995. Multiagent negotiation under time constraints. *Artificial Intelligence* 75:297–345.

Kreps, E., and Wilson, R. 1982. Reputation and imperfect information. *J. of Economic Theory* 27:253–279.

Mas-Colell, A.; Whinston, M.; and Green, J. R. 1995. *Microeconomic Theory.* Oxford University Press.

Osborne, M. J., and Rubinstein, A. 1990. *Bargaining and Markets.* Academic Press, Inc.

Riley, J. 1980. Strong evolutionary equilibrium and the war of attrition. *J. of Theoretical Biology* 82:383–400.

Rosenschein, J. S., and Zlotkin, G. 1994. *Rules of Encounter.* MIT Press.

Roth, A.; Murnighan, J.; and Schoumaker, F. 1988. The deadline effect in bargaining: some experimental evidence. *American Economic Review* 78:806–823.

Rubinstein, A. 1982. Perfect equilibrium in a bargaining model. *Econometrica* 50.

Rubinstein, A. 1988. A bargaining model with incomplete information about time preferences. *Econometrica* 53:1151–1172.

Sandholm, T. W., and Lesser, V. R. 1997. Coalitions among computationally bounded agents. *Artificial Intelligence* 94(1):99–137. Special issue on Economic Principles of Multiagent Systems.

Sandholm, T. W. 1993. An implementation of the contract net protocol based on marginal cost calculations. In *AAAI*, 256–262.

Smith, J. 1974. The theory of games and evolution in animal conflict. *J. of Theoretical Biology* 47:209–221.